

*Compaq StorageWorks™
Virtual Replicator*

System Administrator's Guide

Order number: **AA-RENKB-TE**

Version 1.1

Compaq Computer Corporation
Houston, Texas

June 1999

Possession, use, or copying of the software described in this documentation is authorized only pursuant to a valid written license from Compaq or an authorized sublicensor.

While Compaq believes the information included in this publication is correct as of the date of publication, it is subject to change without notice.

Compaq Computer Corporation makes no representations that the use of its products in the manner described in this publication will not infringe on existing or future patent rights, nor do the descriptions contained in this publication imply the granting of licenses to make, use, or sell equipment or software in accordance with the description.

© 1998,1999 Digital Equipment Corporation

All rights reserved.

Compaq, Compaq logo, DIGITAL, StorageWorks, registered in U.S. Patent and Trade-mark Office.

Intel is a registered trademark of Intel Corporation.

Microsoft, PowerPoint and Windows NT are registered trademarks of Microsoft Corporation.

All other trademarks are the property of their respective holders.

Contents

Preface	ix
1 Introduction	
1.1 Features.....	1-1
1.1.1 Storage Pooling.....	1-2
1.1.2 Virtual Disks.....	1-2
1.1.3 Snapshots.....	1-2
1.1.4 Network Disks.....	1-4
1.1.5 Cluster-Aware.....	1-4
1.1.6 In Action Together.....	1-4
1.2 Benefits.....	1-6
1.3 Creating Snapshots.....	1-8
1.4 Serving Network Disks.....	1-10
1.5 Cluster Resources.....	1-12
1.6 SNMP Support.....	1-13
2 Installation	
2.1 Components.....	2-1
2.2 Prerequisites.....	2-1
2.3 Installing Compaq Storageworks Virtual Replicator.....	2-2
2.3.1 Installing Compaq Storageworks Virtual Replicator Using Microsoft Systems Management Server.....	2-2
2.3.2 Installing Acrobat Reader.....	2-3
3 Planning	
3.1 Pools, Virtual Disks, and Snapshots.....	3-1
3.2 Storage Units in Pools.....	3-2
3.3 Rules for Using Virtual Disks and Snapshots.....	3-3
3.4 Pool Capacity.....	3-3
3.5 Disk Space Requirements for a Virtual Disk.....	3-4
3.5.1 Using the Snapshot Planner.....	3-4

3.6 Disk Space Requirements for Pools	3-5
3.7 Advantages and Disadvantages of Large Pools	3-5
3.8 Existing Backup Tools	3-7
3.8.1 Checking if your Backup Tool Records Information on the Disk.	3-7
3.9 Supported Network Disks	3-8
3.9.1 Disks Serving as Network Disks	3-9
3.9.2 Naming Conventions	3-9
3.9.3 Uniqueness of Names	3-10

4 Guided Tour

4.1 Create a Pool.	4-1
4.2 Create a Virtual Disk in the Pool	4-7
4.3 Create a Snapshot of the Virtual Disk	4-11
4.4 Serve the Snapshot to a Remote Computer	4-14

5 Using the Snap-ins and Commands

5.1 Understanding MMC and Snap-ins	5-1
5.1.1 MMC Panes	5-3
5.1.2 Adding a Snap-In.	5-4
5.2 Using the Command Line Interface	5-6
5.2.1 Using Commands to Manage Remote Computers.	5-8
5.2.2 At the command prompt	5-8
5.2.3 At the Windows NT prompt	5-8
5.2.4 Abbreviating Commands	5-9
5.2.5 Using Uppercase or Lowercase in Commands	5-9
5.3 Security and Privileges	5-9
5.3.1 Managing other domains	5-10

6 Managing Pools

6.1 Creating Pools, Virtual Disks, and Snapshots	6-2
6.1.1 Creating a Pool.	6-2
6.1.2 Creating a Virtual Disk	6-5
6.1.3 Creating a Snapshot	6-6
6.1.4 Scheduling Automatic Snapshots	6-8

6.2 Pool Space	6-9
6.2.1 Monitoring Pool Free Space	6-9
6.2.2 Understanding Delspace	6-11
6.3 Adding Storage Units to a Pool	6-12
6.4 Restoring Virtual Disks	6-15
6.5 Using Snapshots to do Backups	6-16
6.6 Defragmenting Virtual Disks	6-18
6.7 Mapping and Unmapping Drives	6-18
6.7.1 Mapping a Drive Letter to a Virtual Disk or Snapshot	6-18
6.7.2 Partitioning and Formatting a Virtual Disk	6-19
6.7.3 Unmapping a Drive Letter from a Virtual Disk or Snapshot	6-20
6.8 Displaying Information	6-21
6.8.1 Displaying Information about Pools	6-21
6.8.2 Displaying Information about Virtual Disks and Snapshots	6-25
6.9 Deleting Snapshots, Virtual Disks, and Pools	6-27
6.9.1 Deleting a Snapshot	6-27
6.9.2 Deleting a Virtual Disk	6-28
6.9.3 Deleting a Pool	6-29

7 Managing Network Disks

7.1 Serving a Disk	7-1
7.1.1 Serving Virtual Disks and Snapshots	7-2
7.1.2 Serving other Disks	7-3
7.2 Controlling Which Clients Connect to the Disks	7-7
7.3 Connecting to a Served Network Disk	7-8
7.4 Changing Settings	7-12
7.4.1 Changing the IP Subnet Configuration	7-13
7.4.2 Changing the Client System for a Served Network Disk	7-14
7.5 Using a Served Network Disk as the Cluster Quorum Disk	7-14
7.5.1 Steps When the Cluster Already Exists	7-15
7.5.2 Steps If the Cluster Does Not Already Exist	7-16
7.6 Tuning Tips	7-16

7.7 Showing Information	7-17
7.7.1 Showing Information about Served Network Disks	7-17
7.7.2 Showing Information about Connected Network Disks .	7-19
7.8 Disconnecting from a Served Network Disk	7-21
7.9 Stopping the Serving of a Network Disk	7-22

8 Compaq Batch Scheduler Wizards

8.1 Task Automation Wizards	8-1
8.1.1 Using the Batch Scheduler Wizards	8-2
8.1.2 Modifying Task Automation Schedules	8-3

9 Troubleshooting

9.1 Application using a Network Disk Fails or Hangs	9-1
9.2 Lost Delayed Write Errors	9-4
9.3 Applications Return Failed Write Errors	9-5
9.4 Accidentally Delete a Virtual Disk.	9-6
9.5 Snap-in and Commands Do Not Show All the Pools in a Cluster.	9-7
9.6 Serving Tasks Missing from the Snapshot Manager Snap-in	9-7
9.7 Disk Structure Corrupt Errors when Accessing a Snapshot	9-8
9.8 Disk Structure Corrupt Errors when you Stop Serving a Disk.	9-9
9.9 Pool Free Space Fell by More Than the Amount of Data Written	9-10
9.10 Cannot Connect to a Network Disk in Another Subnet . .	9-10
9.11 SCE Connected Disk Resource Fails to Come Online in a Cluster.	9-11
9.12 Reformatting a Virtual Disk Does Not Work.	9-11
9.13 Performance Monitor Shows Zero Network Disks.	9-12
9.14 Reconstructing a Pool	9-12
9.15 Unable to Delete a Snapshot or a Virtual Disk.	9-12

A SnapMgr Commands

DRIVES	A-2
MANAGE.....	A-4
Manage a Remote Computer	A-5
Manage the Local Computer.....	A-6
Show the Managed Computer.....	A-7
POOL	A-8
Create a Pool.....	A-9
Add a Storage Unit to a Pool	A-12
Delete a Pool.....	A-15
Show Pools	A-16
SNAPSHOT	A-21
Create a Snapshot	A-22
Map a Drive Letter to a Snapshot.....	A-24
Unmap the Drive Letter from a Snapshot.....	A-26
Delete a Snapshot	A-27
Show Snapshots	A-29
UNITS.....	A-31
VIRTUALDISK	A-33
Create a Virtual Disk	A-34
Map a Drive Letter to a Virtual Disk	A-36
Partition and Format a Virtual Disk	A-38
Unmap the Drive Letter from a Virtual Disk	A-39
Delete a Virtual Disk	A-40
Show Virtual Disks.....	A-42

B NDMgr Commands

DRIVES	B-2
MANAGE.....	B-4
Manage a Remote Computer	B-5
Manage the Local Computer.....	B-6
Show the Managed Computer.....	B-7
REGISTER	B-8
Add a Registered Client	B-9
Remove a Registered Client	B-11

Show Registered Clients	B-12
SERVE	B-14
Serve a Disk	B-15
Stop Serving a Disk	B-18
Show Served Disks	B-20
USE	B-23
Connect to a Network Disk	B-24
Disconnect from a Network Disk	B-28
Show Connected Network Disks	B-30

Glossary

Index

Preface

Welcome to *Compaq StorageWorks™ Virtual Replicator*.

This book is for system administrators who are experienced with the following:

- Management of Microsoft Windows NT servers and workstations
- Configuration and management of Microsoft Cluster Server (MSCS) clusters
- Cluster Administrator.

Throughout this book, the term **cluster** means an MSCS cluster, and **standalone computer** means a computer that is not in a cluster.

As well as this book, we supply these sources of information:

Where	What
In the box	A <i>Commands at a glance</i> quick reference card.
On the CD-ROM	In the Documentation folder of the SWVR distribution CD-ROM, you will find: <ul style="list-style-type: none">• A readme file that contains the release notes.• Planning charts to help you plan your migration• The PDF version of this book, so you can print extra copies of it
On the web	Visit http://www.compaq.com/storageworks/
Online help	Full information is available in the online help.

1 Introduction

Compaq StorageWorks™ Virtual Replicator (SWVR) provides advanced, centralized storage management capabilities in both clustered and non-clustered Windows NT computing environments. Its innovative storage management features simplify storage configuration and management, and enhance availability and scalability.

This software runs on standalone computers and on computers in Microsoft Cluster Server (MSCS) clusters.

This chapter gives an overview of the product and describes the following topics:

- Features (Page 1-1)
 - Storage Pooling (Page 1-2)
 - Virtual Disks (Page 1-2)
 - Snapshots (Page 1-2)
 - Network Disks (Page 1-4)
 - Cluster-Aware (Page 1-4)
- Benefits (Page 1-6)
- Creating Snapshots (Page 1-8)
- Serving Network Disks (Page 1-10)
- Cluster Resources (Page 1-12)
- SNMP Support (Page 1-13)

1.1 Features

Compaq StorageWorks Virtual Replicator provides manageable storage. It allows you to:

- Group hardware arrays or physical disks to form a large **pool** of storage.
- Divide the pool into **virtual disks** of any size, up to 1 terabyte.
- Make instantaneous copies, called **snapshots**, of the virtual disks and use the snapshots locally or serve them over the network.
- Use the virtual disks on the **local** computer, or **serve** them over the **network** to remote computers.

1.1.1 Storage Pooling

Compaq StorageWorks Virtual Replicator enables the grouping of hardware array storage or physical disks into a logically concatenated pool of disk space. You can create any number of pools, using industry-standard storage components. Any storage to which Windows NT has direct access can be used in a pool. As well as standard single-spindle disks, you can use controller-based fault tolerant disk arrays, such as StorageWorks RAID arrays, referred to as **storage units**.

The storage units provide disk space for the pool in the same way that the physical disks that make up a StorageWorks RAID array do. The *Compaq StorageWorks Virtual Replicator* software controls what data is stored on what virtual disk.

1.1.2 Virtual Disks

The virtual disks that you create in the pool perform and behave just like physical disks. That is, you can format and map drive letters to them, and read from and write to them, just like physical disks.

Disk virtualization allows you to optimally tailor disk space to the size required by users and their applications. You can make the sizes of the virtual disks match the requirements of your applications and users. For example, if a user needs 650 MB of disk space, you can create a 650 MB virtual disk. And if you have a terabyte database, you can combine several StorageWorks RAID arrays in a single pool, and create a terabyte virtual disk that spans the StorageWorks RAID arrays. The size of a virtual disk is limited only by how much free space there is in the pool when you create it.

1.1.3 Snapshots

Compaq StorageWorks Virtual Replicator allows you to make instant replicas (called **snapshots**) of virtual disks in a matter of seconds. Snapshots enable the instantaneous creation of multipurpose virtual

replicas of production data without the requirement of physically copying data. Snapshots function identically to ordinary physical disks with both read and write capability.

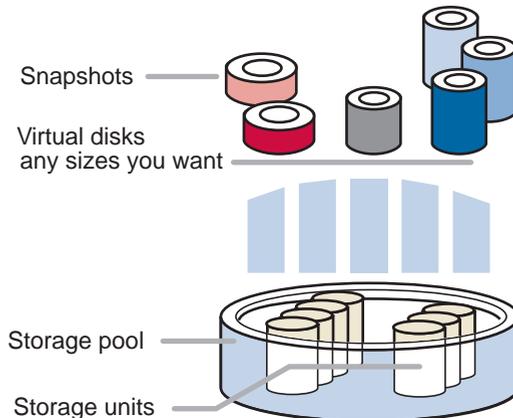
Snapshots are used any time you need a quick copy of your production data without disrupting running applications. You can use snapshots to do online backup and restore, test new applications, or populate your data warehouse or web server. *StorageWorks Virtual Replicator* provides wizards, in conjunction with *Compaq Batch Scheduler*, to automatically schedule unattended creation, deletion, and backup of snapshots. See Chapter 8 for information on these wizards.

A snapshot looks like an exact copy of the original virtual disk, made at an instant in time. It has the same capacity and label, and it contains *exactly* the same data. It is as if you had a camera and took a picture of every single byte of data stored on the original virtual disk at a single instant in time.

When you first create a snapshot it uses no disk space. It starts to use disk space only when you write data to either the original virtual disk or the snapshot itself.

To users and applications, the snapshot looks like a normal disk. You can read from it and write to it, just like a normal disk, and you do not have to interrupt your users to create it. You can create snapshots while users are reading from and writing to the original virtual disk.

The following figure shows a storage pool that is built from two StorageWorks RAID arrays. There are three virtual disks and three snapshots in the pool. The virtual disk on the left has one snapshot, the virtual disk in the middle has no snapshots, and the virtual disk on the right has two snapshots.



1.1.4 Network Disks

In addition to using virtual disks and snapshots directly on your Windows NT servers, *StorageWorks Virtual Replicator* allows disks (virtual, physical, or snapshots) to be served over the network to remote computers.

The *Virtual Replicator* network disk component makes a disk on one computer or cluster behave as if it is physically plugged into another computer or cluster. On the remote computer or cluster, the served disk looks just like a locally attached disk. Applications do not know that it is located elsewhere on the network. This capability makes it possible to centralize storage where it can be physically secure and regularly backed up, while the remote network-based clients continue to operate as though all their storage was locally connected. You can install any Windows NT applications on the served disk and can use it with applications that need local disks and do not work with shares.

1.1.5 Cluster-Aware

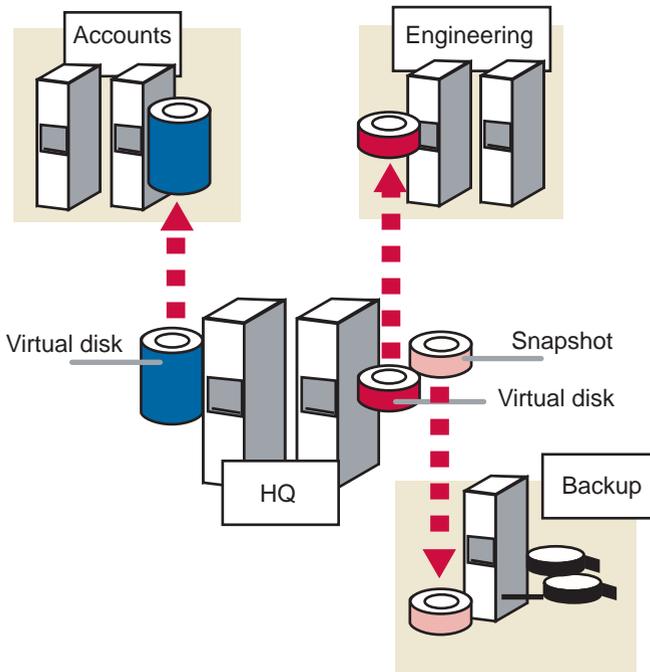
Compaq StorageWorks Virtual Replicator runs on top of MSCS when cluster support is required, to provide higher availability for data and applications. Failover and failback of pools, virtual disks, and snapshots are performed as a unit in an MSCS cluster. Cluster resources are automatically created to simplify cluster management.

1.1.6 In Action Together

You can make a virtual disk available in a variety of ways:

- Windows NT File Service
- DFS mount point
- *Virtual Replicator* network disk

In the following figure, a data center has a cluster called HQ that provides disk storage for clusters in the accounts and engineering departments. The HQ cluster makes one virtual disk available to each department.



The data center also has a dedicated backup server. In the figure, it is backing up a snapshot of the engineering department's virtual disk. The engineers continue reading from and writing to their disk while the system manager creates the snapshot and backs it up.

When the accounts department asks for another 500 MB of disk space, the system administrator does not have to go to the accounts department and add another physical disk to the Accounts cluster. The system administrator uses the *Compaq StorageWorks Virtual Replicator* point-and-click interface to create a new 500 MB virtual disk on the HQ cluster, and make it available to the Accounts cluster.

The storage components described in the previous section are fully cluster-aware. For example, in the previous figure, if the configuration is using network disks and if the node in the HQ cluster that is currently serving the virtual disk to the Accounts cluster fails or is taken offline, the pool that contains the virtual disk automatically fails over to the other node in the HQ cluster, and that node automatically serves the virtual disk to the Accounts cluster.

Similarly, if the node in the Accounts cluster that is currently connected to the virtual disk fails, the connection automatically fails over to the other node in the Accounts cluster.

1.2 Benefits

These are some ways that typical enterprises can benefit from *Compaq StorageWorks Virtual Replicator*.

■ Optimize disk storage by using storage pooling/virtualization

System administrators can optimize disk storage by centralizing multiple physical disks and hardware StorageWorks RAID arrays to form a logically concatenated pool of storage.

By binding the centralized disks into pools of storage, you can partition the individual pools into multiple virtual disks. The virtual disks are available to users either on the local server or networked to remote servers when and where they are needed using network disks.

■ Use virtual and network disks to allocate disk space flexibly

When you buy a new disk, there is no need to dedicate it to one server. You can add it to an existing pool or use it in a new pool, and divide it into virtual disks of whatever sizes you want.

For example, if you buy a 100 GB StorageWorks RAID array, you can add it to an existing pool. You can do this online, while users are reading from and writing to the virtual disks and snapshots in the pool.

Partition the disk space and serve it to your application servers like this:

- Create a 5 GB virtual disk for your Microsoft Exchange server.
- Create a 3 GB virtual disk for your web server.

The remaining 2 GB is spare space for your snapshots to use. More information on snapshots and how they use disk space is provided later in this manual.

- **Use snapshots to back up your data online**

To back up a virtual disk online, just create a snapshot of the disk and then back up the snapshot. The files that were open in the virtual disk are automatically closed in the snapshot, so your backup captures the whole disk.

Creating the snapshot takes a matter of seconds, and can be done online, while your users are reading from and writing to the original virtual disk.

Alternatively, if your backup strategy involves quiescing or shutting down your applications to capture all the recent updates that they are buffering in memory, you can restart your applications almost immediately – just as soon as you have created the snapshot.

- **Use snapshots to keep copies of your data online**

It is easy to keep snapshots online, so that when users accidentally delete files, they can copy the files back from the snapshots themselves.

For example, you do daily backups using snapshots, and keep each snapshot online for a week before you delete it. On Thursday, when a user accidentally deletes a file, he or she can copy it from Wednesday's snapshot, instead of asking you to restore the file from backup tapes.

In another example, at the end of business each day, you run a batch job that processes the data in your database to generate reports. Before you start the batch job, you create a snapshot of the virtual disk that holds the database data, so that if the batch job fails part way through, you can run it against the original data again.

- **Use snapshots to test applications**

You can easily and quickly test applications against a snapshot of your live production data without disrupting your business and without any risk of corrupting your live production data.

You cut out the step of having to make a physical copy of your entire production data, and you do not have to take your data offline. You do not have to wait for hours while the data is laboriously copied to spare disks, or waste money buying extra disks just for the full copy of your test data.

- **Use snapshots with data mining applications**

Data mining applications can populate your data warehouse by processing a snapshot of your live production data. You cut out the step of having to make a physical copy of your production data. And you do not have to take your data offline. The data mining applications can extract data from the snapshot without interrupting your business.

- Use snapshots to do backups whenever you want

If you want to save the state of a disk at midday, you do not have to physically copy the data to tape at midday. Just create a snapshot at midday, then back up the snapshot to tape whenever you want. If there are not any free tape drives or your system is heavily loaded at midday, you can delay the tape copy operation until a tape drive becomes available or the load on your system is lighter.

- Use network disks to locate disks where you want

Each Windows NT server needs its own startup (boot) disk, but you can decide where to locate its other disks. They do not have to be physically attached to it. You can centralize your disk storage instead of having it dispersed throughout your organization. Centralized storage is easier and cheaper to manage.

You can choose to locate the disks in a single central location, or you may prefer to have one location for each geographical area.

1.3 Creating Snapshots

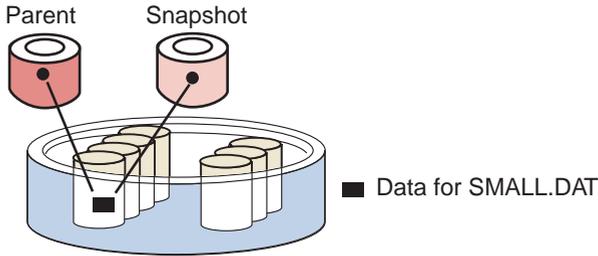
When you create a snapshot, the disk of which you are making a virtual replica is called the **parent** disk. You can create more snapshots from the same parent disk, or you can create snapshots of the snapshots.

The original virtual disk, all of its snapshots, and all the snapshots of those snapshots, and so on, are known as a **family**.

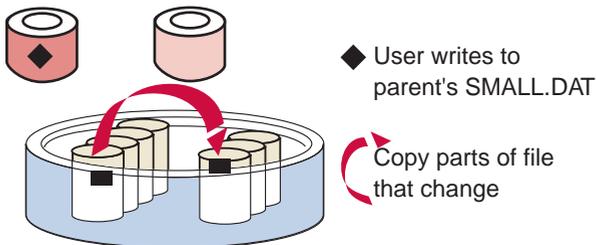
Creating a snapshot is fast – a matter of seconds – because it does not involve *copying* any data. When you first create the snapshot, it holds exactly *the same* data as its parent disk, so there is no need to make a physical copy of the data.

The snapshot does not use any disk space. It shares the same physical disk space as its parent. But if users modify either the snapshot or its parent, the data being modified must be physically copied beforehand. A snapshot starts to use disk space only when the parent or snapshot is modified.

For example, when a snapshot is first created, it has a file called SMALL.DAT which is identical to the file called SMALL.DAT on the parent disk. So both files share the same disk space, as shown in the following figure.

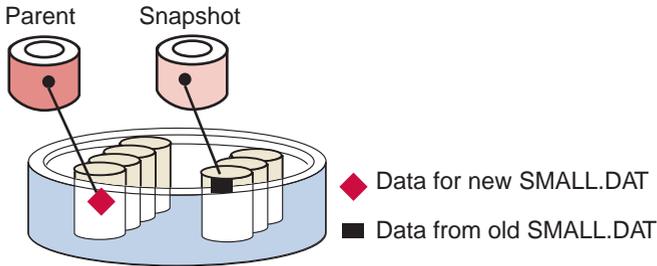


A user then writes to the parent's SMALL.DAT. Before that write happens, the SWVR software automatically makes a separate physical copy for the snapshot of the *parts* of the file that are changed. This operation is called a **copy-out**. The software then updates the parent disk with the new data.



In our example, the file is very small so the whole file is copied, but for a normal file, only the parts that changed would be copied.

The snapshot's SMALL.DAT now contains the old data, and the parent's SMALL.DAT contains the new data. All this is done automatically, by the SWVR software.



A snapshot starts to consume disk space when you create, modify or delete the files or directories that are stored on either the snapshot or the parent disk.

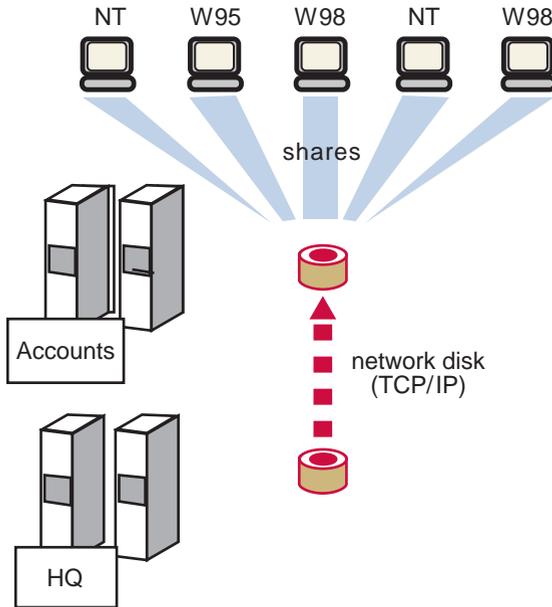
The snapshot consumes disk space in units called **segments**, which are 32 KB by default. Each copy-out operation consumes a segment (32 KB) of disk space in the pool. The segment size is set when you create the pool.

When a segment has been copied-out, subsequent writes to the segment do not incur copy-out operations. The snapshot has its own copy of the segment, so there is no need to copy it again.

1.4 Serving Network Disks

Compaq StorageWorks Virtual Replicator provides the optional feature of serving disks over the network, to be used with any disk, including virtual disks and snapshots.

In the following figure, the HQ cluster is serving a disk over the network to the Accounts cluster. On the Accounts cluster, the network disk looks like a locally attached disk. Just as with a normal locally attached disk, users in the Accounts department use shares to access the data stored on the network disk.



Network disks are similar to shares, except:

- With a share, you share **part** of a disk – you share a folder on the disk. With a network disk you serve the **whole** disk.
- With a share, one computer **shares** part of its disk with other computers. With a network disk, one computer or cluster **gives** its disk to another computer or cluster. The computer or cluster serving the disk cannot even access the disk directly – the disk is marked as offline on that computer or cluster.
- With a share, lots of computers can connect to the share. With a network disk, only one standalone computer or cluster can connect to the network disk.
- With a share, although you map a drive letter to it, like a local disk, the share does not look like a normal local disk. You cannot install some applications directly onto it, you cannot create shares on it, and some applications cannot store their data on it.

But a network disk looks exactly like a local disk. You can install any application directly onto it, you can create shares on it, and any application can store its data on a network disk. You can also use it as a cluster quorum disk.

The performance of network disks is similar to that of local disks, provided that your network is fast enough and is not already heavily loaded. This is because the design is simple and involves minimal overhead – you are only serving raw disk blocks. Requests to read and write raw disk blocks are simply transferred over the network using TCP/IP.

You can serve normal physical disks, virtual disks, or snapshots. You can serve disks over a TCP/IP network.

When there are two network links between the network disk client and server, if one link fails, the software automatically uses the other link. Setting up a network disk is a two step process:

Step 1	You serve the disk. The computer serving the disk has a list of registered clients . This is a list of all the computers that can connect to its served disks. You do not specify valid clients.
Step 2	The client computer or cluster connects to the served disk. The attempt to connect fails if it is not on the list of registered clients on the computer serving the disk.

Step 1 is the equivalent of creating a share. Step 2 is the equivalent of mapping a network drive to a share, which connects the client computer to the share. In Step 2, a drive letter is automatically mapped to the network disk on the client computer.

1.5 Cluster Resources

Compaq StorageWorks Virtual Replicator provides four cluster resources:

Resource type	Description
SCE Storage Unit	A storage unit in a pool
SCE Pool	The pool itself
SCE Served Disk	A network disk that is being served by the cluster
SCE Connected Disk	A network disk that the cluster is connected to

- ▲ **Tip:** The virtual disks and snapshots in a pool automatically appear on whichever node currently owns the pool resource. Therefore, there is no need for a resource for virtual disks or snapshots.

In a cluster, when you create pools and when you serve or connect to network disks, *Virtual Replicator* automatically creates cluster resources for you, so that your pools and network disks can fail over within the

cluster. This saves you from having to use Cluster Administrator to set up the resources.

The *Virtual Replicator* software automatically creates the necessary groups and resources, and sets up the correct dependencies, possible owners, and so on.

For example, when you create a pool, the *Virtual Replicator* software creates a new cluster group for the pool. This group contains the pool resource and a storage unit resource for each unit in the pool. The pool resource is dependent on all of the storage unit resources. A pool must always be able to access all of its storage units, therefore, all of the storage units must be online before the pool resource comes online.

And when you serve a virtual disk in the pool, the *Virtual Replicator* software automatically creates a served disk resource in the pool group, and makes it dependent on the pool resource.

If the node that currently owns the pool resource fails or is taken offline, the pool's group fails over to another node in the cluster, which then serves the virtual disk. Users on the client (receiving) computer that is connected to the served disk continue to access the disk, unaware that it is now being served by a different node in the cluster. No I/Os are lost, and there is no interruption to the service.

1.6 SNMP Support

Compaq StorageWorks Virtual Replicator supports the use of Simple Network Management Protocol (SNMP) for the exchange of management information between management console applications such as UniCenter and managed entities. When you install *Virtual Replicator*, you also install the following:

- **SNMP agent**

A processing element that retrieves and updates local management information based on requests from applications that are monitoring and controlling the nodes.

- **Management Information Base (MIB)**

A collection of managed objects in a database, which defines the variables in the tables, the data to be retrieved, and the format in which to present the data.

The SNMP agent implements the actual returning of the data to whichever node asks for it. The MIB is installed wherever you installed *Virtual Replicator* on the local machine. The default location is:

\Program Files\Compaq\SWVR\

Locate the MIB and copy it to the node where your management console applications reside.

For more information on using your SNMP management application, refer to your SNMP documentation.

In addition to SNMP support, *Compaq StorageWorks Virtual Replicator* also supports Microsoft Windows NT Performance Monitor. The Performance Monitor shows the regular Counter values for virtual disks, as well as snapshots. The counters are similar to those for physical disks. If you go to the Performance Monitor, select Add to Chart, select logical disk, and select the drive (virtual disk or snapshot); you can choose any counter that is valid for a logical disk.

2 Installation

This chapter describes how to install the following products:

- *Compaq StorageWorks Virtual Replicator*
 - Storage Software
 - Snapshot Planner
- Adobe Acrobat Reader

2.1 Components

Version 1.1 of *Compaq StorageWorks Virtual Replicator* provides two separately installable components:

- Storage Software

This component allows you to create pools of storage, divide the pools into virtual disks of whatever size you want, and create virtual copies, called snapshots, of the virtual disks. It also allows you to serve disks over the network. You can serve any disks, including virtual disks and snapshots.
- Snapshot Planner

This component is a planning tool that helps you predict how much disk space you will need for your snapshots. You can use it before you install the Storage Software component.

2.2 Prerequisites

To install *Compaq StorageWorks Virtual Replicator* V1.1, you must have the following software installed:

- Microsoft Windows NT Version 4.0, with Service Pack 3 or greater
- Microsoft Internet Explorer™ Version 4.0 or greater

If you do not have Internet Explorer, you can download a free copy of the latest version of Internet Explorer from the Microsoft web site at:

<http://www.microsoft.com/>

2.3 Installing *Compaq Storageworks Virtual Replicator*

Follow these steps to install SWVR:

1. Log on to a user account in the **local Administrators** group.
2. Insert the SWVR installation CD, select the component you want to install, and then follow the instructions on your screen.
3. The next time you start your computer, make sure you log on to the **same user account** as in Step 2.
4. In a cluster, if you installed the Storage Software component, install **exactly the same** Storage Software option on the other node(s) in the cluster. The Storage Software component is fully cluster aware, so you must install the same Storage Software option on all nodes of the cluster. For example, if you installed the network disk client software on one node, make sure you install the network disk client software on the other nodes too.

2.3.1 Installing *Compaq Storageworks Virtual Replicator* Using Microsoft Systems Management Server

Microsoft Systems Management Server (SMS) allows you to install *Compaq StorageWorks Virtual Replicator* as an SMS package. You can use the package definition file, SWVR.PDF, to define the package parameters. SWVR.PDF is located on the installation CD in the following directory:

`\Storage Software\swvr.pdf`

This directory also contains the file, setup.iss, which provides installation commands that inform the SWVR installation process which components to install.

By default, an SMS installation of *StorageWorks Virtual Replicator* installs with the following options:

- Create pools, virtual disks, and snapshots
- Serve network disks over the network to remote computers (Network disk server).

- Manage *StorageWorks Virtual Replicator* on the local and remote computers.

If you need to install the Network disk client software, copy the contents of the Storage Software folder from the CD-ROM to a hard drive. Then, copy the file \Storage Software\client.iss to \Storage Software\setup.iss, and create a new SMS package.

To return the SMS installation to the default installation options, copy the file \Storage Software\server.iss to \Storage Software\setup.iss, and create a new SMS package.

- ▲ Do not delete or rename the files server.iss and client.iss. Copying server.iss or client.iss to setup.iss overwrites the previous contents of setup.iss, this enabling SMS to install *StorageWorks Virtual Replicator* with a different set of installation options.

For information on how to create packages for Microsoft Systems Management Server, see the documentation for SMS.

2.3.2 Installing Acrobat Reader

A PDF version of this Manual is included in the Documentation folder on the installation CD. In the same folder, you will find an Adobe Acrobat Reader kit, in case you do not have a PDF viewer.

The Adobe Acrobat Reader kit is an Intel kit, so if you have an Alpha, you need to install FX!32 to use it. You will find the FX!32 kit in the same directory on the installation CD.

3 Planning

This chapter describes what information you need to consider when planning the configuration. The topics described in this chapter are:

- Pools, virtual disks, and snapshots (Page 3-1)
- Storage units in a pool (Page 3-2)
- Rules for using virtual disks and snapshots (Page 3-3)
- Pool capacity (Page 3-3)
- Disk space requirements for virtual disks (Page 3-4)
- Disk space requirements for pools (Page 3-5)
- Advantages and disadvantages of large pools (Page 3-5)
- Existing backup tools (Page 3-7)
- Supported network disks (Page 3-8)
 - ▲ To help plan the *Compaq StorageWorks™ Virtual Replicator* configurations, use the planning chart in the Documentation folder on the SWVR distribution CD-ROM. There is a blank chart to fill in and an example of a filled-in chart.

3.1 Pools, Virtual Disks, and Snapshots

Compaq StorageWorks Virtual Replicator allows you to have any number of pools on each standalone computer or cluster. The number is limited only by the number of storage units that are available. The following numbers are allowed for virtual disks and snapshots:

- Up to 8 virtual disks in each pool.

The capacity of each virtual disk must be at least 10 MB, and is limited only by the amount of free space in the pool.

- Up to 12 snapshots in each family.

For example, you can have 12 snapshots of each virtual disk, or you can have 11 snapshots of a virtual disk, and one of these snapshots can have a snapshot.

3.2 Storage Units in Pools

When deciding which storage units to use in a pool, note the following:

- Create pools from whole disks, not from partitions or logical drives on a disk.
- Each storage unit must be a raw, unformatted disk that does not contain any partitions.
- A pool can contain up to eight storage units. They can have different capacities and be from different manufacturers.
- Do not use network disks that are being served by remote computers or clusters.
- In a cluster, use disks only on the shared storage bus; do not use local disks.
- Do not use removable disks, such as floppy disks or Jaz drives.
- Use standard single-spindle disks or controller-based fault-tolerant disk arrays.

When you create a virtual disk or snapshot in a pool, you cannot specify which storage units it uses. The virtual disks and snapshots that you create in the pool can use disk space from anywhere in the pool.

Compaq recommends that all the storage units in a pool have the same redundancy, read-write, and failure characteristics. For example, they should all be StorageWorks RAID 5 arrays, or they should all be mirror arrays, or they should all be standard disks. Mixing different types of storage units in the same pool could result in unpredictable characteristics.

If you want the data on a virtual disk to have particular characteristics, create it in a pool whose storage units all have those characteristics.

For example, if you want the data on a virtual disk to be mirrored, create the virtual disk in a pool whose storage units are all controller-based mirror arrays. Note that you must not use Disk Administrator to create a mirror set on a virtual disk (see “Rules for Using Virtual Disks and Snapshots” on Page 3-3).

Another reason why you should make sure that all the storage units in a pool have the same characteristic is that a pool is a single point of failure. This means that if one of its storage units becomes inaccessible due to a hardware fault, you lose the whole pool and all of its virtual disks and snapshots.

If this happens, you have to reconstruct the pool. You have to create a brand new pool, then create new virtual disks in the pool, and then restore data from your backup tapes to the new virtual disks.

For example, there is no point in using a standard single-spindle disk and a StorageWorks RAID 5 array in the same pool. Although the StorageWorks RAID array can survive the failure of any of the individual disk spindles that make up the set, you lose the whole pool if the standard disk fails. The pool is as weak as its weakest storage unit, in this case, the standard disk.

3.3 Rules for Using Virtual Disks and Snapshots

You can use a virtual disk or snapshot just like a normal disk, except:

- You cannot use it as your system disk.

- You cannot create more than one partition on it.

Partitioning is a way of dividing up a big disk. With virtual disks, you just create disks of whatever sizes you want; there is no need to partition them.

- The partition must be formatted with the NTFS file system.

Like the MSCS cluster software, *Compaq StorageWorks Virtual Replicator* does not support the FAT file system.

- Do not use Disk Administrator to create a volume set, or a mirror or stripe set on it.

Like the MSCS cluster software, *Compaq StorageWorks Virtual Replicator* does not support software RAID.

3.4 Pool Capacity

The capacity of a pool is slightly less than the combined capacities of its storage units, because the pool configuration data uses some disk space.

The amount of space used by the configuration data varies and can be up to 10% of the combined capacities of the storage units.

For example, if you create a pool from five 20 GB StorageWorks RAID arrays, for planning purposes, assume that the capacity of the pool is 90

GB. You can increase the capacity of a pool at any time by adding storage units to it. You can do this dynamically, while users are reading from and writing to the virtual disks and snapshots in the pool.

Note that you cannot remove storage units from a pool in this version of *Compaq StorageWorks Virtual Replicator*.

3.5 Disk Space Requirements for a Virtual Disk

The disk space that a virtual disk uses is the same as the capacity of the virtual disk. When you create a virtual disk, you specify its capacity in megabytes, and the free space in the pool drops by that number of megabytes.

The disk space a snapshot uses can vary from 0 MB to its own capacity, depending on the rate at which your data changes, and for how long you keep the snapshot.

When you first create a snapshot, it does not use any space. It consumes space only when users modify the data stored on either the snapshot itself or its parent. The more data they modify, the more space the snapshot uses. At worst, the snapshot could use as much space as its own capacity (which is the same as the capacity of its parent disk).

If you have been taking incremental backups of your data, you will have a good idea of how rapidly your data changes. You can use the size of your incremental save arrays as a guide. A snapshot will probably take up less space than this because the incremental save set saves a whole file even though only a part of the file changed.

3.5.1 Using the Snapshot Planner

Compaq StorageWorks Virtual Replicator provides a tool called Snapshot Planner to help you predict how much space your snapshots will use.

Snapshot Planner tracks actual reads and writes to your existing drives and works out how much space a snapshot of each drive would need and how many extra I/Os the snapshot would cause due to copy-out operations.

Snapshot Planner is very useful if you plan to migrate all the data on an existing hard disk drive to a virtual disk of the same capacity. It tracks the I/Os to the existing drive, and lets you work out exactly how much extra space you would need for its snapshots.

To run Snapshot Planner, click **Start**, then **Programs**, then **Compaq StorageWorks Virtual Replicator**, and then **Snapshot Planner**.

For more information, see the Snapshot Planner's online help.

- ▲ **Tip:** Use Snapshot Planner **before** you set up your pools and virtual disks. Snapshot Planner cannot see virtual disks, so once you have created virtual disks in a pool, you cannot use it to predict the cost of snapshots of your virtual disks.

3.6 Disk Space Requirements for Pools

Just like physical disks, you must never let your pools run out of disk space. If a pool runs out of disk space, writes to the pool are failed. If the applications using the disks in the pool do not use the system cache, you will realize that the pool is full immediately, the applications immediately report failed writes.

Most applications use the system cache to buffer writes to disk. In this case, you may not realize that the pool is full until some time later, when the cache tries to write the data to disk. The write from the cache fails and you see a popup window telling you that a delayed write has failed and you may have lost some data.

When a pool runs out of disk space, you must either add a storage unit to it, or delete one of its snapshots or virtual disks.

- ▲ **Tip:** If a pool is full, do not try to free up space by deleting file, because deleting files does not free up any space in the pool.

If a pool does not contain any snapshots, it can never run out of disk space. But if a pool contains snapshots, and its capacity is less than the combined capacities of all of its virtual disks and snapshots, there is a possibility that it could run out of disk space and become full. In this case, monitor its free space regularly and make sure that it does not become full. For more information, see "Pool Space" on Page 6-9.

3.7 Advantages and Disadvantages of Large Pools

When deciding whether to create a large pool that contains lots of virtual disks or lots of small pools, each containing fewer virtual disks, note the following:

- Decide which applications you want in each pool, so you can give each application its own virtual disk, instead of having them sharing virtual disks.

- In a cluster, a pool is a unit of failover. When the node that the pool is on fails, the pool and all the applications using its virtual disks fail over to another node in the cluster.

To manually balance the load and control which applications run on which node, create small pools. If you have one pool for each application, you have maximum control.

- Create a large pool to use storage units that provide redundancy to protect against hardware faults, such as StorageWorks RAID 5 arrays. This avoids the problem of a storage unit becoming inaccessible due to a hardware fault, and the loss of the whole pool and all the data in it.
- To balance the load across individual storage units, create single storage unit pools. *Compaq StorageWorks Virtual Replicator* does not do any load balancing across the different storage units in a pool.
- By creating large pools that contain several virtual disks, you can reduce the cost of disk space for snapshots.

For example, assume that you want to create three 10 GB virtual disks. You want to use snapshots only to back up the disks, so the only time that any of these disks have a snapshot is when it is being backed up. To guarantee that you do not run out of disk space during a backup, assume that the snapshot needs as much space as its own capacity, namely 10 GB.

If you can arrange to back up the three disks one at a time, you can create a single 40 GB pool for all of them. You create three 10 GB virtual disks in the pool, leaving 10 GB of spare capacity for the snapshot of whichever disk is being backed up.

This saves 50% more disk space than creating a pool for each disk, which would need three 20 GB pools, because the virtual disks “time-share” the 10 GB of spare capacity for their snapshots.

Look at which storage units you have available and work out how much disk space to give to each application. If you have lots of small disks and an application that needs a large amount of disk space, you may decide to create one pool out of the small disks.

Alternatively, if you have lots of applications that need a small amount of disk space and you have only a big disk, you might want to make a pool out of the one big disk and divide it into lots of small virtual disks.

Maybe you do not want to use snapshots and are interested in saving disk space. You have 5 applications that each need 10 GB of disk space, but you have only 20 GB disks. You decide to combine 3 of the disks to make a single pool. The capacity of the pool will be about 54 GB (about 10% less

than the combined capacities of the storage units), enough for five 10 GB virtual disks.

If you want to use snapshots, you need to take into account the disk space that the snapshots will use.

For example, you have a cluster that is running an Exchange mail server and an SQL database:

- Exchange needs a 16 GB disk, and you want to use snapshots only to back it up.
Snapshot Planner tells you that the snapshot needs 2 GB of disk space. You create a 20 GB pool for Exchange (that is 2 extra GBs, just to be on the safe side).
- The data in your 20 GB SQL database changes rapidly. Snapshot Planner tells you that the snapshot needs 10 GB of disk space. You create a 35 GB pool (that is 5 extra GBs, just to be on the safe side).

3.8 Existing Backup Tools

You can use your existing backup tools to back up virtual disks and snapshots. If you want to use snapshots to do backups **and** you want to do incremental or differential backups, check whether your backup tool records information about the backup on the disk being backed up.

If your backup tool records information about the backup on the disk being backed up, always do full backups when you use snapshots to do backups. You must not do incremental or differential backups, because you delete the snapshot when you have finished the backup. When you delete the snapshot, you lose the record of the backup, and so subsequent incremental and differential backups do not work properly.

3.8.1 Checking if your Backup Tool Records Information on the Disk

Some backup tools record information about the backup by clearing the archive bit of the files that are backed up. To check whether your backup tool does this, in Windows NT Explorer look at the Properties sheet for your files before and after a backup. If your backup tool clears the archive bit, before the backup the archive attribute box is checked for files that have been modified since the last backup, and after the backup it is no longer checked.

Another way to check if your backup tool records information on the disk being backed up is as follows:

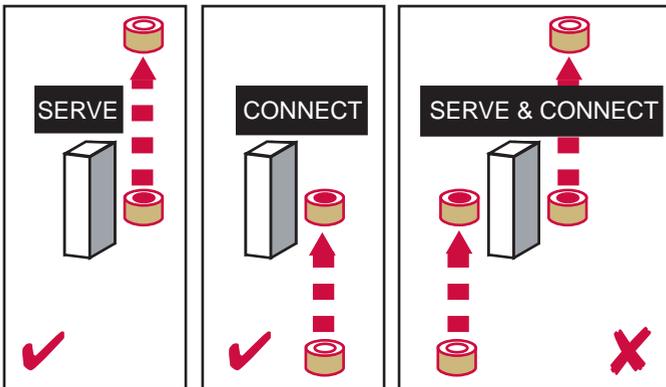
1. Create a snapshot of a virtual disk, then do a full backup of the snapshot disk.
2. Delete the snapshot.
3. Edit a few files on the virtual disk, then create a new snapshot. Do an incremental backup of the new snapshot.
4. If the backup save sets produced by the full and incremental backups in Steps 1 and 3 are the same size, your backup tool is recording information about the backup on the disk being backed up. Your incremental backup saved every file, not just the ones that had changed.
5. Delete the snapshot.

For instructions on how to use snapshots to do backups, see “Using Snapshots to do Backups” on Page 6-16.

3.9 Supported Network Disks

When planning your network disk configurations, bear in mind these points:

- In this release of *Compaq StorageWorks Virtual Replicator*, the network disk server software cannot run on the same computer as the network disk client software. This means that a single computer cannot both serve and connect to network disks.



- In a cluster, you must **not** configure one node to serve disks and the other to connect to served disks. Do not install the network disk server software on one node and the network disk client software on another node.

- Each standalone computer or cluster can serve or connect to up to 23 served network disks (drive letters D to Z).
- The standalone computer or cluster serving a network disk must have the same architecture as the standalone computer or cluster that connects to it. They must both be Intel or both be Alpha.
For example, an Alpha computer should not serve network disks to an Intel computer.

3.9.1 Disks Serving as Network Disks

When planning which disks you want to serve as network disks, bear in mind these points:

- You cannot serve removable disks, such as floppy disks or Jaz drives.
- A network disk must have only one partition.
Partitioning is a way of dividing up a big disk. With *Compaq StorageWorks Virtual Replicator*, you can use pools and virtual disks to divide your available disk space; you do not need to use partitions.
- The partition must be formatted with the NTFS file system.
Like the MSCS cluster software, *Compaq StorageWorks Virtual Replicator* does not support the FAT file system.
- Do not use Disk Administrator to create a volume set, mirror set or stripe set on the disk.
Like the MSCS cluster software, *Compaq StorageWorks Virtual Replicator* does not support the Windows NT fault tolerant (FT) driver.

The client computers that connect to your network disks must not repartition them or use Disk Administrator to create volume, mirror, or stripe sets on them.

3.9.2 Naming Conventions

You choose a name for every pool, virtual disk and snapshot you create. And when you serve a disk, you have to choose a network disk name, that the remote client specifies when connecting to it.

- ▲ **Tip:** Take time to plan what names you want to use, because you cannot change them later.

The names can be up to 23 characters and are not case-sensitive. So Mypool_1 is the same as myPOOL_1.

Compaq recommends that you follow these naming conventions:

- When you serve a virtual disk or snapshot, make its network disk name the same as its virtual disk or snapshot name. This is the default when you use the *Virtual Replicator* graphical interfaces.
- If you want to create snapshots of snapshots and keep track of the hierarchy within the family, choose snapshot names that indicate the hierarchy, such as SnapOfA and SnapOfSnapOfA. The SWVR software shows which family each snapshot is in, but it does not show the parent relationship for snapshots of snapshots.

3.9.3 Uniqueness of Names

Every pool, virtual disk and snapshot must have a different name. If you have a pool called Pool1, you cannot create a snapshot called Pool1.

In a cluster, no two cluster groups can have the same name, and no two cluster resources can have the same name. Because *Virtual Replicator* automatically creates cluster groups and resources using the names you choose, there are additional restrictions on names in clusters:

- A pool cannot have the same name as any existing cluster group or cluster resource. So if there is a cluster resource called Res1, you cannot create a pool called Res1.
- If you serve a virtual disk or snapshot as a network disk, the name of the virtual disk or snapshot must be different from that of any existing cluster resource. So if there is a cluster resource called Cluster1, do not call a virtual disk Cluster1 if you want to subsequently serve the virtual disk.
- If you serve a disk other than a virtual disk or snapshot, its network disk name must be different from that of any pool or any existing cluster group or cluster resource. So if there is a cluster group called Group1, you cannot call the network disk Group1.

4 Guided Tour

This chapter provides a guided tour that gives an overview of the *Compaq StorageWorks™ Virtual Replicator* tools that manage pools and network disks. This tour performs the following tasks:

- Creates a Pool (Page 4-1)
- Creates a Virtual Disk in the Pool (Page 4-7)
- Creates a Snapshot of the Virtual Disk (Page 4-11)
- Serves the Snapshot to a Remote Computer (Page 4-14)

It should take you about 15 minutes to read through the tour – you do not need to follow the steps on your own computer.

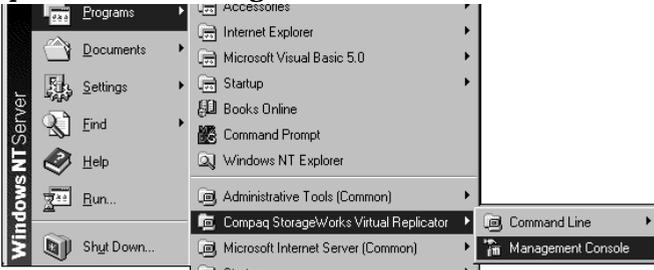
4.1 Create a Pool

This part of the tour creates a pool on the local HQ cluster using 2 disks on the cluster's shared SCSI bus as the storage units for the pool.

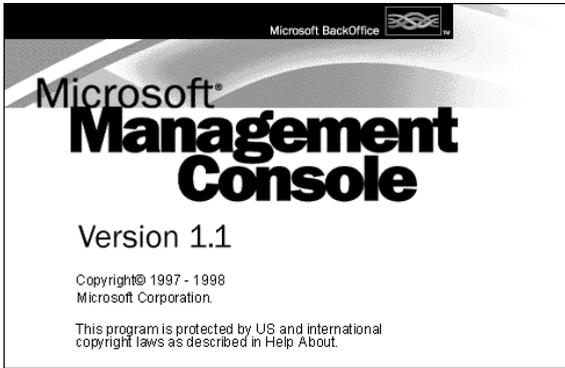
You create pools out of **raw** disks, therefore, assume that the disks do not contain any partitions and that they are not configured as cluster resources of type Physical Disk. See Section 6.1 for more information on creating pools.

Also, assume that you have already installed the snapshot and network disk server components of *Virtual Replicator* on both nodes of the HQ cluster.

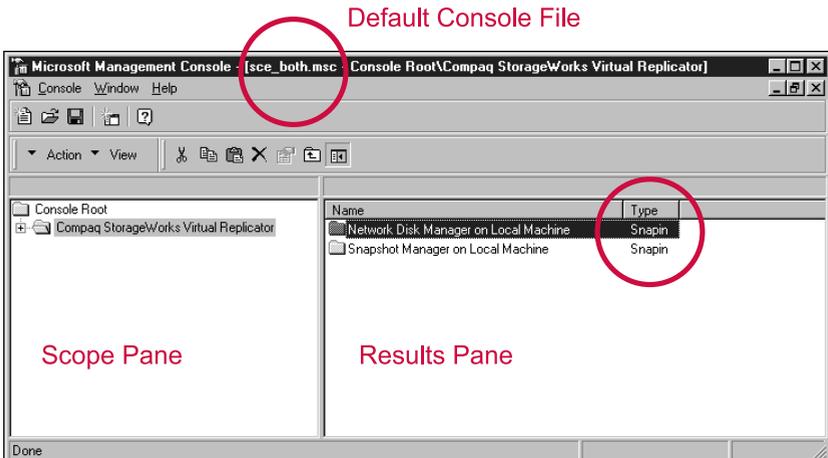
1. Click **Start**, then **Programs**, then **Compaq StorageWorks Virtual Replicator**, then **Management Console**.



The Microsoft Management Console (MMC) splash screen is briefly displayed.



And then this screen is displayed.



MMC is started and you have opened the default **console file**, which is provided to help get you up and running quickly.

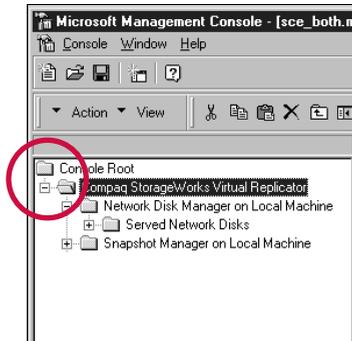
The default console file depends on which components you installed. In this example, it is called `sce_both.msc` and has two **snap-ins** – the MMC term for a tool that is set up to manage a particular computer or cluster.

- ▲ **Tip:** In a cluster, the default snap-ins are labelled Network Disk Manager on Local Machine and Snapshot Manager on Local Machine, but they actually manage the local cluster, not just a single node in the cluster.

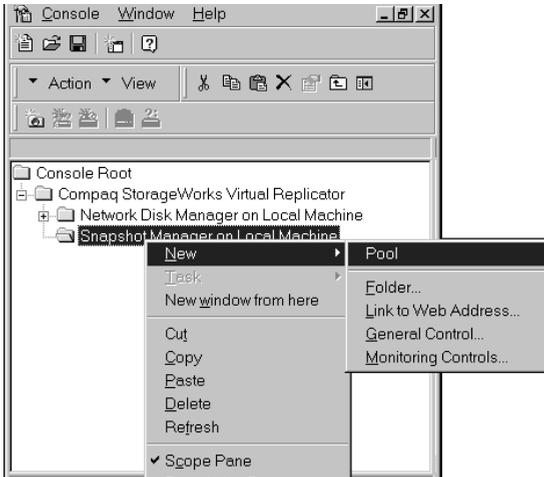
MMC is like Windows NT Explorer. In the left-hand **Scope** pane, you use folders to organize your snap-ins into a tree view, in the same way that you use folders on a disk to organize your files.

You click on things in the left-hand Scope pane to expand them. And the right-hand **Results** pane displays information about whatever you have selected in the left-hand Scope pane.

2. Click the plus sign to the left of the **Compaq StorageWorks Virtual Replicator** folder, and open the folders that appear under it.



3. Select **Snapshot Manager on Local Machine**, and then right-click **New** then **Pool**.

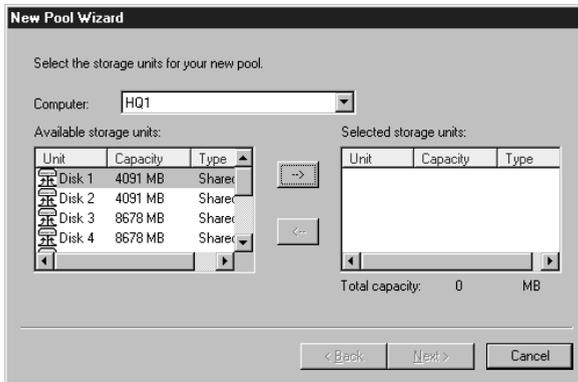


- ▲ **Tip:** In MMC, all the options available on the right-click menus are also available on the **Action** menu. For example, in this step, instead of using the right-click menu, we could have clicked **Action**, and then **New** then **Pool**. These actions are also available as buttons on the toolbar.

4. The New Pool wizard welcome screen is displayed.



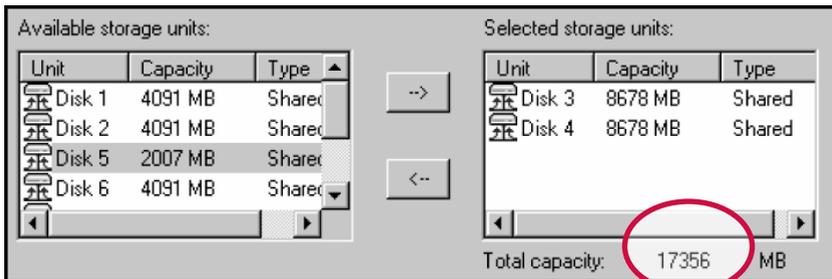
5. Click **Next**. The New Pool Wizard screen is displayed.



On the left, there is a list of available storage units. The Type column tells you whether the disk is shared (on the cluster's shared storage bus) or local. Only select shared disks in a cluster.

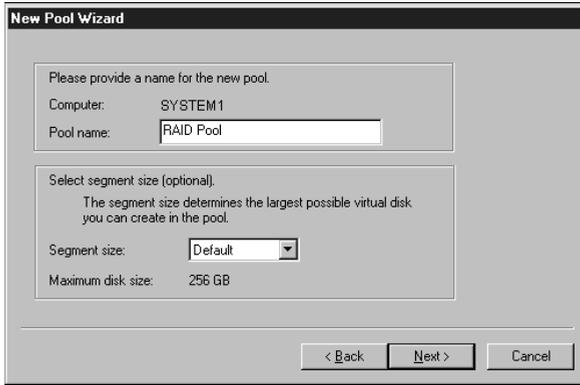
6. Select Disk 3 and Disk 4, then click Right Arrow.

The disks move to the right side of the screen.



The total capacity field shows the sum of the capacities of the disks you selected. The capacity of the pool could be up to 10% less than this, because some of the disk space is used by internal pool configuration data.

7. Click **Next**. On the next screen, type the name that you want to give to the new pool. This example calls it RAID Pool. Then supply the



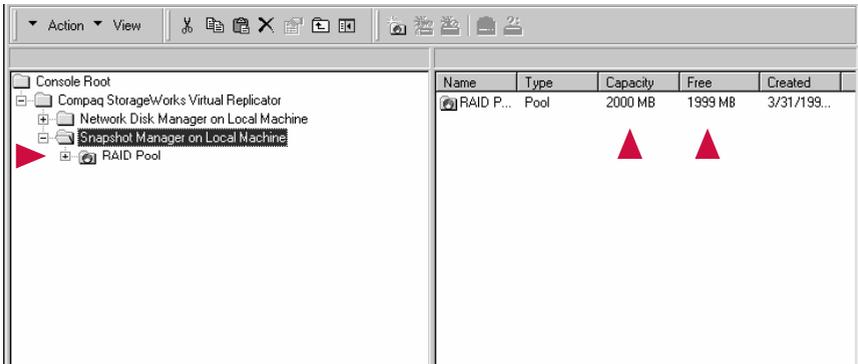
segment size. The default is 32K.

8. Click **Next**. The final wizard screen is displayed.



This screen tells us whether you created the pool. Like all the *Virtual Replicator* final wizard screens, if the wizard fails, it explains why it failed on this screen.

9. Click **Finish** to dismiss the final wizard screen. Now the management console looks like this.

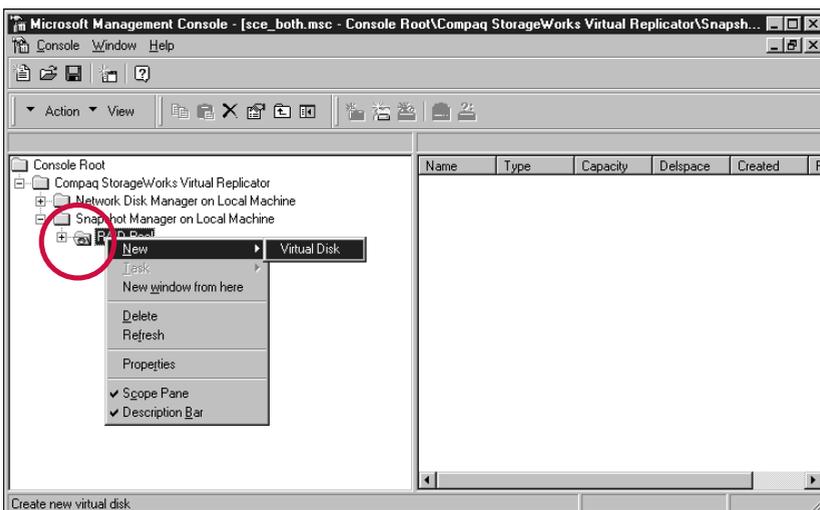


The new pool is listed. Notice that the pool's capacity is 1 MB more than its free space. That is because the capacity is rounded up and the free space is rounded down.

4.2 Create a Virtual Disk in the Pool

This part of the tour creates a 1500 MB virtual disk, called HQ Data, in the pool. You will map a drive letter to the virtual disk, format it, and copy some files onto it.

1. Select RAID Pool, and then right-click **New** then **Virtual Disk**.
2. The welcome screen of the New Virtual Disk wizard is displayed. Click **Next** to start the wizard.



3. On the next screen, type the name of the new virtual disk and its capacity – This example calls it HQ Data and specifies 1500 MB – then click **Next**.

New Virtual Disk Wizard

Please provide a name and capacity for the new virtual disk. Its capacity cannot exceed the free space in the pool.

Pool

Name:	RAID Pool
Free space:	17287

New Virtual Disk

Name:	HQ Data
Capacity:	1500 MB

< Back Next > Cancel

4. On the next screen, select the drive letter that you want to map to the virtual disk. Select V for virtual, then click **Next**.

New Virtual Disk Wizard

You must map a drive letter to the virtual disk before you can format it.

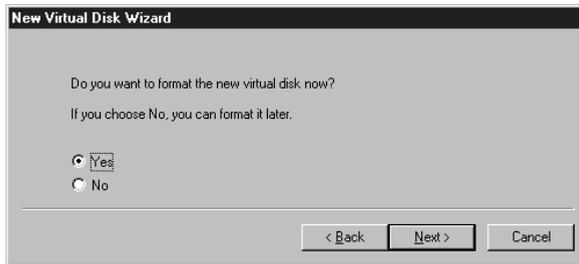
Do you want to map a drive letter now? If you choose not to, you can map a drive letter later.

Map drive letter V

No drive letter

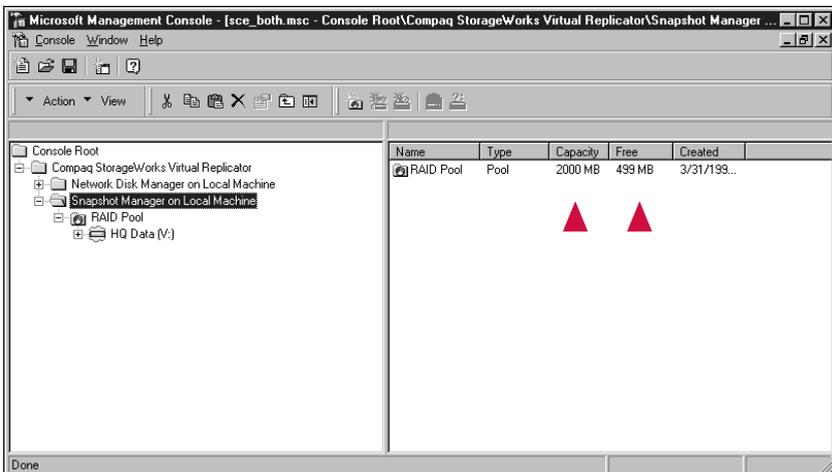
< Back Next > Cancel

- Click **Next** to format the virtual disk.



This creates a single partition on the disk, formats it with the NTFS file system, and gives it the same label as the virtual disk name (HQ Data).

- The final wizard screen is displayed. Check that there are no error messages, then click **Finish** to dismiss the screen.
- Select the **Snapshot Manager on Local Machine** folder like this.



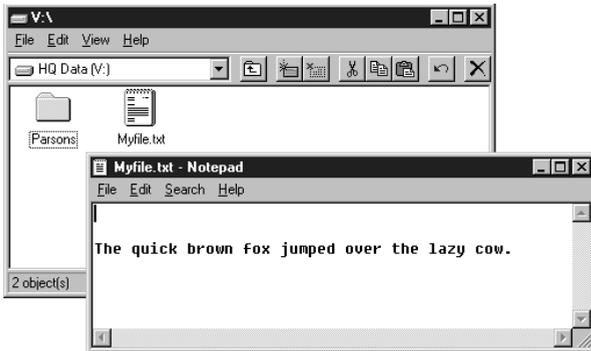
On the left, under RAID Pool, the new virtual disk, HQ Data, is indicated. The drive letter you mapped to it, V, is in brackets. On the right, notice that the pool's free space is now 499 MB. It has dropped by 1501 MB, the capacity of the virtual disk you have just created.

8. The new disk appears in My Computer:

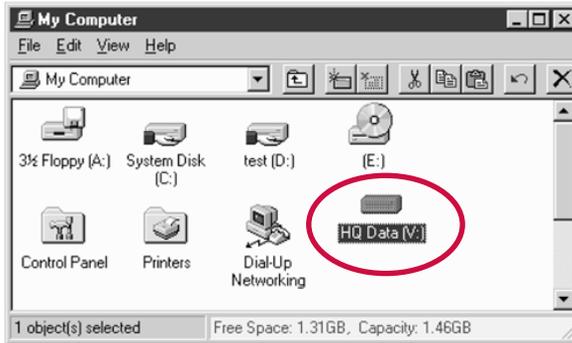


It is drive letter V, volume label HQ Data.

9. Copy a folder called Parsons and a file called Myfile.txt to the new disk:



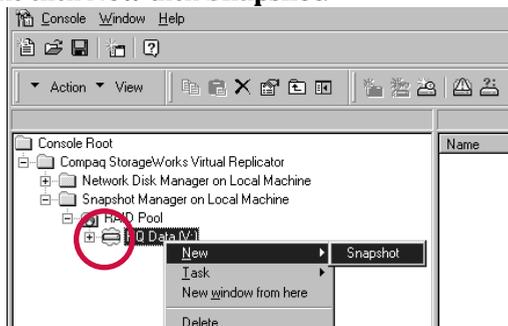
10. Select the disk in My Computer, and look at the bottom status bar. The disk has a capacity of 1.46 GB (0.04 GB is being used by the file system for internal configuration data) and 1.31 GB of free space:



4.3 Create a Snapshot of the Virtual Disk

This part of the tour creates a snapshot, called HQ Data Snapshot, of the virtual disk HQ Data. You will map a drive letter to the snapshot and check that it contains the same files as its parent disk. Then, you will check the free space in the pool to find out how much disk space the snapshot is using.

1. In MMC select the virtual disk, HQ Data, that you just created, and then right-click **New** then **Snapshot**.

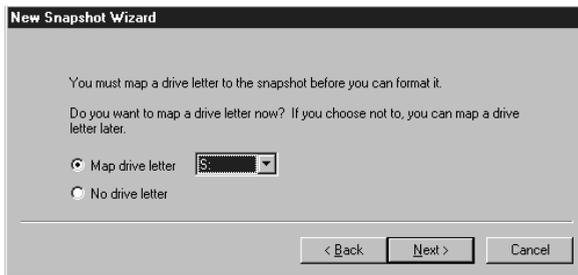


2. The welcome screen of the New Snapshot wizard is displayed. Click **Next** to start the wizard.

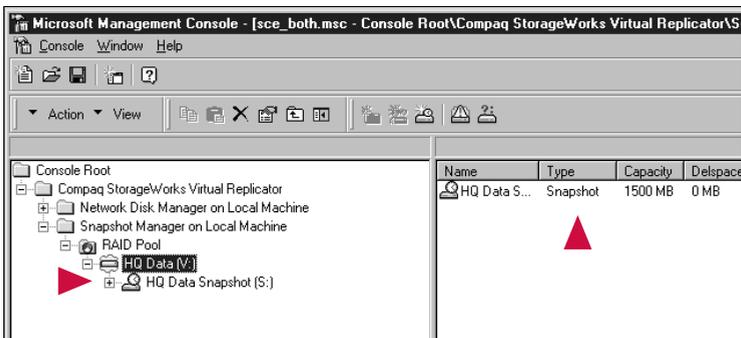
3. Type in the name that you want to give to the new snapshot – This example calls it HQ Data Snapshot – then click **Next**.



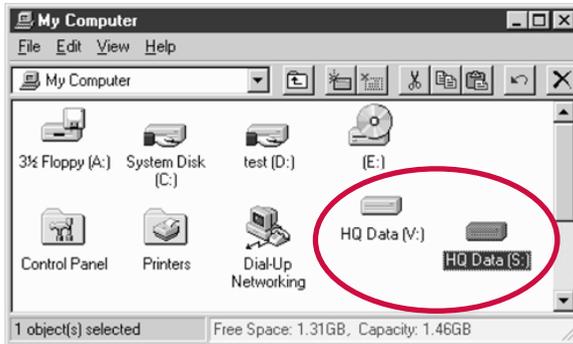
4. On the next screen, click **Map drive letter**, then select the drive letter you want to map to the new snapshot – choose S for snapshot – then click **Next**.



5. The final wizard screen is displayed. Check that there are no error messages, then click **Finish** to dismiss the screen.
6. In the management console, the new snapshot is listed under its parent.

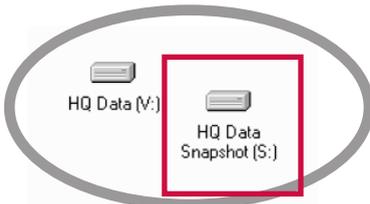


7. Look at My Computer.

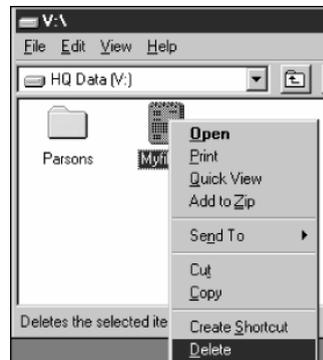


There is the new snapshot – it has the same label (HQ Data) as its parent disk because it is a virtual copy of its parent. Select it and look at its free space and capacity on the bottom status bar. It has the same free space as its parent because it contains exactly the same files.

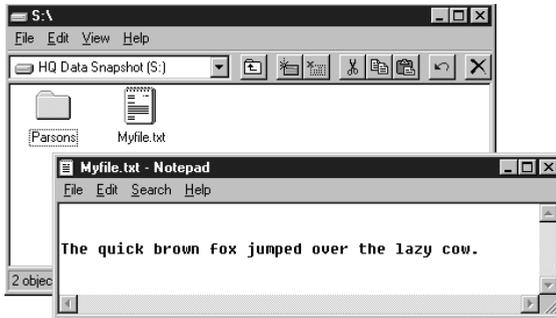
8. Change the label of the snapshot so there is no confusion between the two disks. In My Computer, select the snapshot (drive S), then right-click **Properties** and change its label to HQ Data Snapshot. Click **OK**.



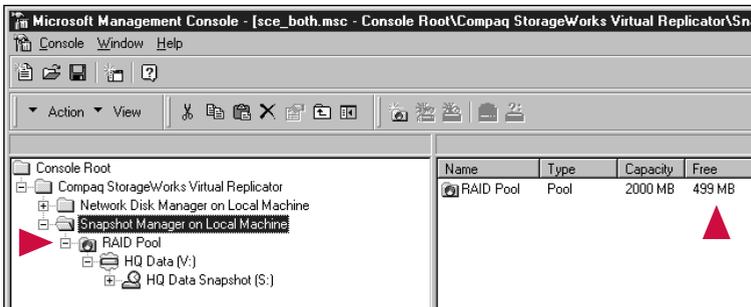
9. Now delete the file Myfile.txt on the parent disk, drive V.



10. But the file is still there on the snapshot disk (drive S).



11. Finally, go back to MMC, and check the pool's free space. It is still 499 MB – the snapshot is not using any disk space.



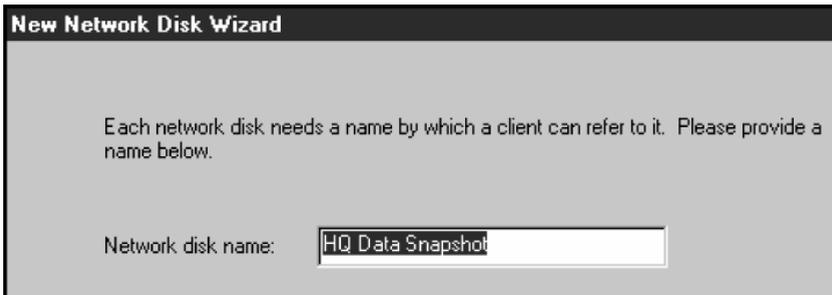
4.4 Serve the Snapshot to a Remote Computer

The final part of the guided tour serves the snapshot to the remote computer named Backup. This involves three tasks:

- Serving the snapshot on the local HQ cluster.
- Adding a snap-in to the MMC console, so that you can manage the remote computer Backup.
- Using this snap-in to connect the Backup computer to the served snapshot.

1. In MMC, select the snapshot HQ Data Snapshot, and then right-click **Task** then **Serve as Network Disk**.

2. The welcome screen of the New Network Disk wizard is displayed. Click **Next** to start the wizard.
3. On the next screen, click **Next** again to accept the default network disk name (which is the same as the snapshot name, HQ Data Snapshot).

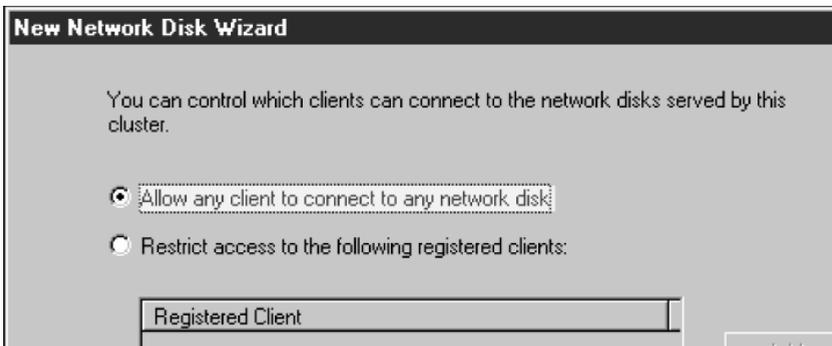


New Network Disk Wizard

Each network disk needs a name by which a client can refer to it. Please provide a name below.

Network disk name:

4. Click **Next** again to accept the default security settings, which allow any client computer to access our network disks.



New Network Disk Wizard

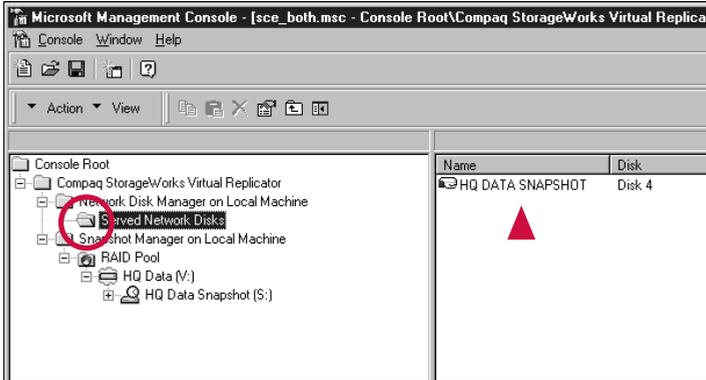
You can control which clients can connect to the network disks served by this cluster.

Allow any client to connect to any network disk

Restrict access to the following registered clients:

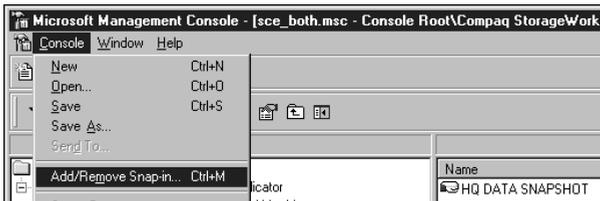
5. The final wizard screen is displayed. Check that there are no error messages, then click **Finish** to dismiss the screen.

- Now select the Served Network Disks folder. The management console shows that HQ Data Snapshot is being served as a network disk.

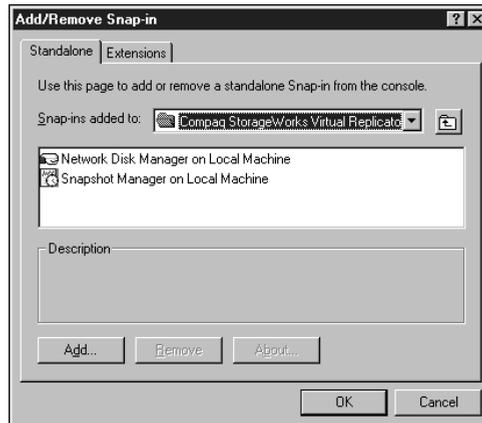


Notice that the name is in uppercase. Network Disk Manager displays all names in uppercase.

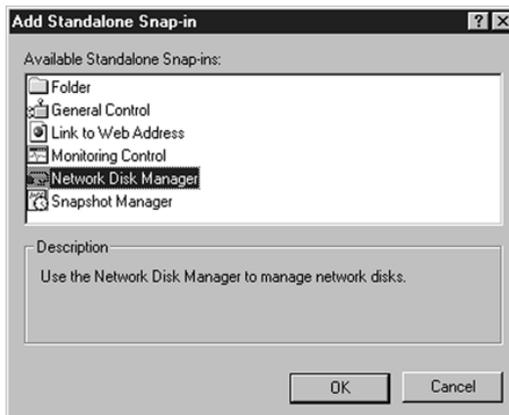
- Now add an MMC snap-in so that you can manage the remote computer Backup. From the MMC **Console** menu, select **Add/Remove Snap-in**.



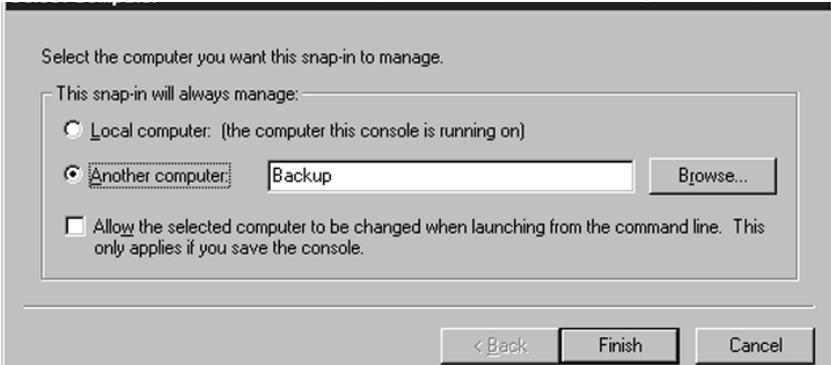
8. Make sure that **Compaq StorageWorks Virtual Replicator** is selected in the top list box, then click **Add**.



9. On the Add Standalone Snap-in screen, select **Network Disk Manager**, then click **OK**.



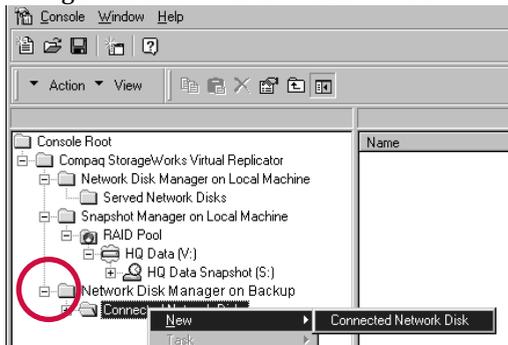
10. On the next screen, click the **Another computer** button, then type in the name of the remote computer you want to manage, Backup.



Click **Finish**, then on the next screen click **OK**. You have added a Network Disk Manager snap-in to manage Backup remotely.

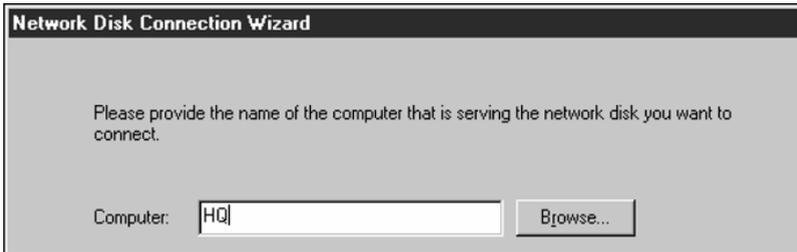
11. Now use the snap-in to connect Backup to the served snapshot.

In the left pane, click the plus sign in front of **Network Disk Manager on Backup** to see the **Connected Network Disks** folder. Select the folder, then right-click **New** then **Connected Network Disk**:

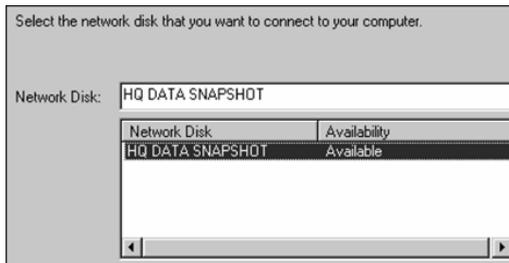


12. The welcome screen of the Network Disk Connection wizard is displayed. Click **Next** to start the wizard.

13. On the next screen, either type the name of the computer or cluster that is serving the network disk or select it from the browse list. Type HQ, and then click **Next**.

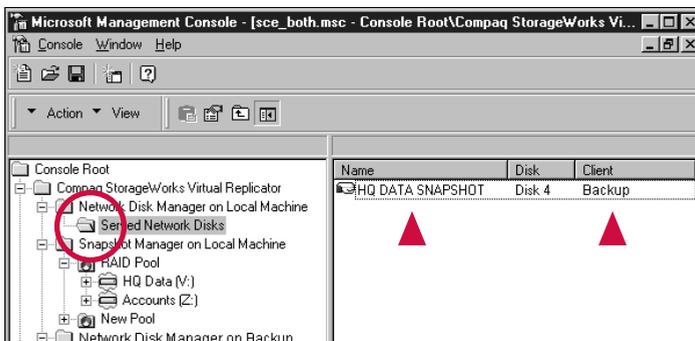


14. On the next screen, select the network disk to connect to on the HQ cluster. Select **HQ DATA SNAPSHOT**, then click **Next**.

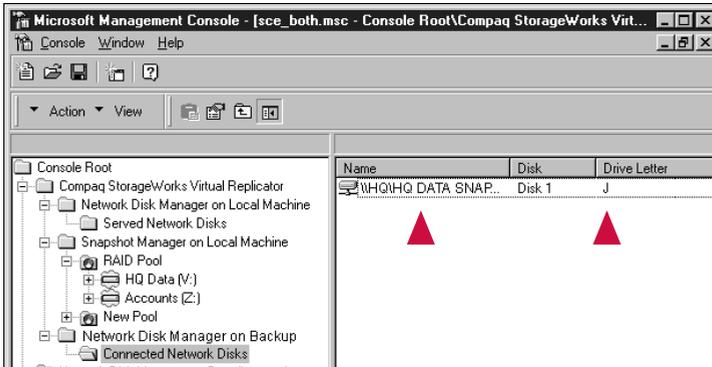


15. The final wizard screen is displayed. Check that there are no error messages, then click **Finish** to dismiss the screen.

16. Now the management console shows that the client **BACKUP** is connected to **HQ DATA SNAPSHOT** which is being served by the local cluster.



17. Select the **Connected Network Disks** folder under the snap-in **Network Disk Manager on Backup**.



The Name column on the right tells you that the network disk called HQ DATA SNAPHOT is being served by HQ.

The Drive Letter column tells you that the drive letter J has been temporarily mapped to the disk on Backup. This was the first available drive letter on Backup. It is not persistent, so the next time Backup restarts a different drive letter may be mapped to the disk.

5 Using the Snap-ins and Commands

This chapter describes how to use the MMC snap-ins and the command-line interface to manage pools and network disks. The following topics are described in this chapter:

- Understanding MMC and Snap-Ins (Page 5-1)
- Adding a Snap-In (Page 5-4)
- Using the Command-Line Interface (Page 5-6)
- Using Commands to Manage Remote Computers (Page 5-8)
- Abbreviating Commands (Page 5-9)
- Using Uppercase and Lowercase in Commands (Page 5-9)
- Security and Privileges (Page 5-9)

The Cluster Administrator software is on the MSCS distribution CD-ROM. Run the MSCS setup program to install the Cluster Administrator software.

- ▲ **Tip:** In a cluster, make sure that you install the same *Virtual Replicator* storage components on every node in the cluster. For example, select the snapshots option on every node, or select the network disk client option on every node.
- ▲ To use a standalone computer to manage pools and network disks on a remote cluster, install the Cluster Administrator software on the standalone computer, as well as the *Virtual Replicator* management tools.

5.1 Understanding MMC and Snap-ins

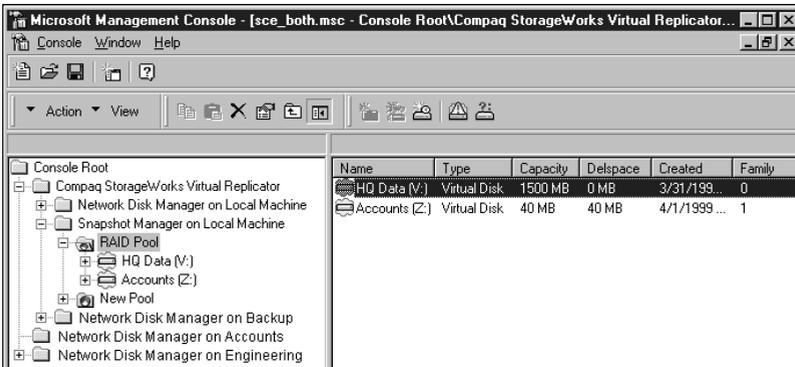
Microsoft Management Console (MMC) is the management environment from Microsoft. It lets you use different management tools, and manage the local and remote computers, all from within the same window on your

desktop. Microsoft and other software vendors are currently migrating their existing management tools into the MMC environment. Eventually, you will be able to perform all of your management tasks from within the same environment.

The management tools that you run in the MMC environment are called **snap-ins**. *Compaq StorageWorks™ Virtual Replicator* provides two snap-ins.

Snap-in	Description
Snapshot Manager	Manages pools, and the virtual disks and snapshots in them. It has a Network Disk Manager extension that lets you use it to serve virtual disks and snapshots.
Network Disk Manager	Manages network disks.

MMC looks and works a bit like Windows NT Explorer. The window is divided into two panes, as shown in this screen capture.



In the left-hand **Scope** pane, you use folders to organize your snap-ins into a tree view, in the same way that you use folders on a disk to organize your files. For example, you can create a folder called Susan that contains all the snap-ins that Susan uses. Or you can organize your snap-ins according to which software vendor supplied them, and have, say a Compaq folder and a Microsoft folder.

You can have many different tree views, and store each one in its own **console file**. For example, each operator could have their own console file, that contains all the snap-ins that they use, arranged in a tree view that suits their particular needs.

Console files have the .msc extension. **Compaq StorageWorks Virtual Replicator** provides a default console file to get you started quickly. This file contains snap-ins to manage the local standalone computer or cluster.

- ▲ **Tip:** In a cluster, the default snap-ins are labelled Snapshot Manager on Local Machine and Network Disk Manager on Local Machine, but they actually manage the local cluster, not just a single node in the cluster.

If you want to manage another standalone computer or cluster, you need to add a snap-in for that computer or cluster. For example, if you want to manage network disks on another cluster, add a Network Disk Manager snap-in and specify the name of the cluster you want to manage, as described in “Adding a Snap-In” on Page 5-4.

5.1.1 MMC Panes

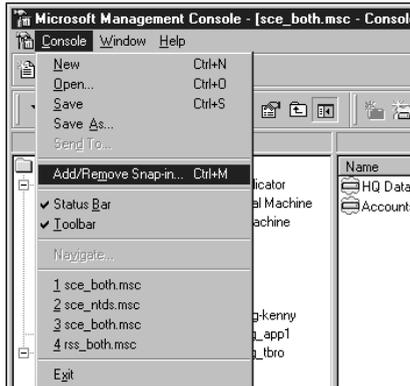
When you select something in the left-hand Scope pane, the results of your selection are displayed in the right-hand **Results** pane. For example, when you select the Snapshot Manager snap-in in the left-hand Scope pane, the right-hand Results pane displays the pools on the standalone computer or cluster that the selected snap-in is managing.

When you select this in the left-hand Scope pane	The right-hand Results pane shows
A Snapshot Manager snap-in	All the pools that are currently online on the computer or cluster that the snap-in is managing.
A pool	All the virtual disks in the pool.
A virtual disk	All the snapshots in the virtual disk's family.
A snapshot	Nothing (the Results pane is empty).
A Network Disk Manager snap-in	Either the Served Network Disks folder or the Connected Network Disks folder.
A Served Network Disks folder	All the network disks that are currently being served by the computer or cluster that the snap-in is managing.
A Connected Network Disks folder	All the network disks that the computer or cluster to which the snap-in is managing is currently connected.

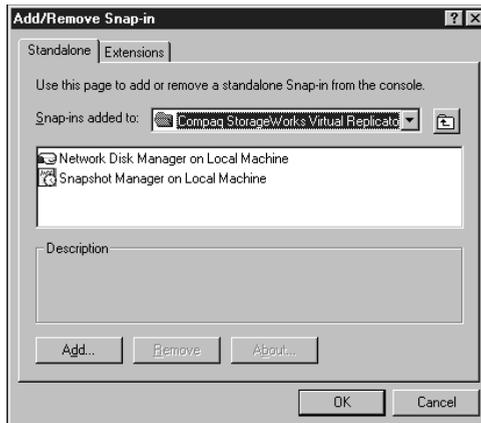
5.1.2 Adding a Snap-In

To add a snap-in to your MMC console:

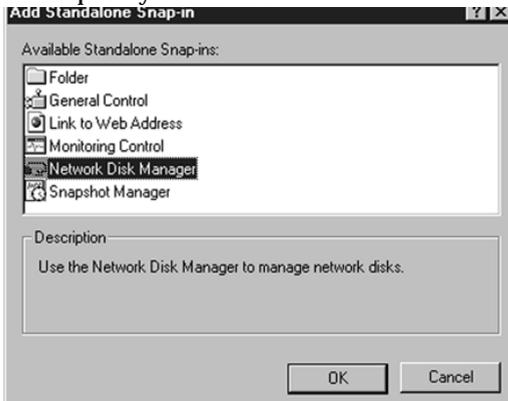
1. On MMC's **Console** menu, click **Add/Remove Snap-in**.



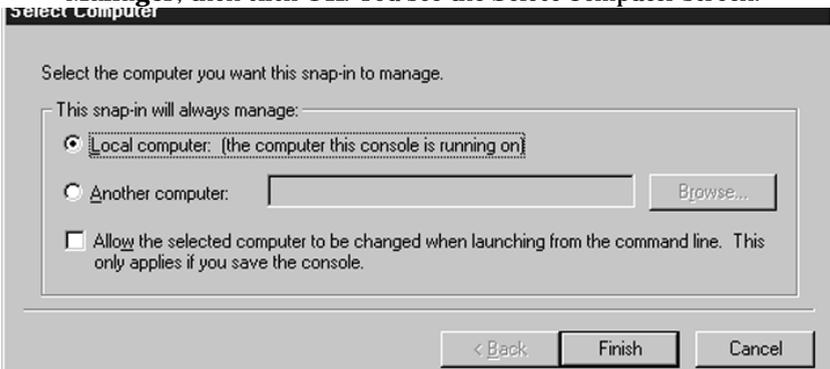
2. On the Add/Remove Snap-in screen, use the top list box to select the folder to which you want to add the snap-in. For example, select the **Compaq StorageWorks Virtual Replicator** folder.



3. Click **Add**. You see the Add Standalone Snap-in screen, which shows all the snap-ins you can add.



4. Select the snap-in that you want to add, for example select **Snapshot Manager**, then click **OK**. You see the Select Computer screen.



- ▲ **Tip:** The square check box controls what happens when you launch the MMC program from the MMC command line. See the MMC documentation for more information.
5. If you want the snap-in to manage the local standalone computer or cluster, click **Finish**.

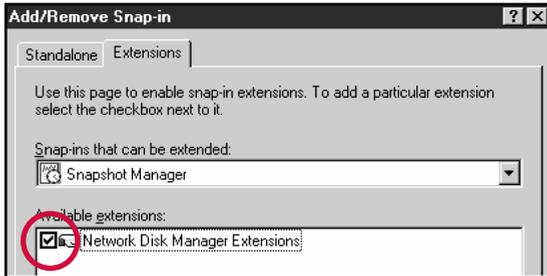
If you want the snap-in to manage another standalone computer or cluster, click the **Another computer** button. Then either type the name of the computer or cluster, or browse for it on the network, and then click **Finish**.

- ▲ **Tip:** If you type or select the name of a node in a cluster, the cluster name is used instead of the computer name. For example, if you are adding a Snapshot Manager snap-in and you type ACCOUNTS1, the snap-in will be labelled

Snapshot Manager on ACCOUNTS, where ACCOUNTS is the name of the cluster in which ACCOUNTS1 resides.

6. If you are adding a Snapshot Manager snap-in **and** you want it to be able to serve virtual disks and snapshots, select the Network Disk Manager extension snap-in.

To do this, click the **Extensions** tab, and check the box to the left of **Network Disk Manager Extensions**, then click **OK**.



5.2 Using the Command Line Interface

Compaq StorageWorks Virtual Replicator provides a command line interface so that you can write batch jobs to automate tasks. For example, you can write batch jobs to back up your data.

There are two sets of commands, which are described in the appendixes of this book. They provide the same functionality as the two snap-ins.

- Use SnapMgr commands to manage pools, and the virtual disks and snapshots in them. They are described in Appendix A.

- Use NDMgr commands to manage network disks. They are described in Appendix B.

SnapMgr command	Description
MANAGE	Controls which computer or cluster you manage (only available at the SnapMgr or NDMgr prompt).
UNITS	Shows which storage units you can use to create a new pool or add to an existing pool.
POOL	Manages pools.
VIRTUALDISK	Manages virtual disks.
SNAPSHOT	Manages snapshots.
DRIVES	Shows which drive letters are available for you to map to virtual disks, snapshots, and the network disks to which you connect.

NDMgr command	Description
MANAGE	Controls which computer or cluster you manage (only available at the SnapMgr or NDMgr prompt).
REGISTER	Manages the list of client computers that are allowed to connect to the network disks that you serve.
SERVE	Manages the network disks that you serve.
USE	Manages the network disks to which you connect.
DRIVES	Shows which drive letters are available for you to map to virtual disks, snapshots, and the network disks to which you connect.

Each set of commands has its own command prefix: SNAPMGR or NDMGR. For example, this command deletes the snapshot called My Snapshot:

```
C:\> SNAPMGR SNAPSHOT "My Snapshot" /DELETE
```

The commands also have their own prompt (SnapMgr or NDMgr). When you issue commands at this prompt, you can omit the command prefix. For example:

```
C:\> SNAPMGR
SnapMgr> SNAPSHOT "My Snapshot" /DELETE
```

- ▲ **Tip:** To create a Windows NT command window that displays the command prompt, click **Start**, then **Programs**, then **Compaq StorageWorks Virtual**

Replicator, then **Command Line**, and then either **Snapshot Manager** or **Network Disk Manager**, depending on which set of commands you want to use.

5.2.1 Using Commands to Manage Remote Computers

Just as with the snap-ins, you can use the commands to manage both local and remote computers and clusters.

By default, you manage the local standalone computer or cluster. The previous examples delete a snapshot on the local computer or cluster.

The way that you manage a remote computer or cluster depends on whether you are at the command prompt or the Windows NT prompt.

5.2.2 At the command prompt

If you are at the command prompt, use the **MANAGE** command to control which computer or cluster you manage.

For example, these commands show which drive letters are available, first on the Accounts cluster, and then on the local cluster:

```
SnapMgr> MANAGE Accounts  
SnapMgr> DRIVES
```

```
Available drive letters: GHKMNPQRST
```

```
SnapMgr> MANAGE /LOCAL  
SnapMgr> DRIVES
```

```
Available drive letters: DGHMNPQ
```

When you have used a **MANAGE** command to manage a remote computer or cluster, use **MANAGE /LOCAL** to revert to managing the local computer or cluster.

5.2.3 At the Windows NT prompt

If you are at the Windows NT prompt, each command automatically manages the local computer or cluster unless you specify the name of a remote computer or cluster after the **SNAPMGR** or **NDMGR** command prefix.

For example, these commands show which drive letters are available, first on the Accounts cluster, and then on the local cluster:

```
C:\> SNAPMGR Accounts DRIVES
```

```
Available drive letters: GHKMNPQRST
```

```
C:\> SNAPMGR DRIVES
```

```
Available drive letters: DGHMNPQ
```

5.2.4 Abbreviating Commands

You can abbreviate the commands to three letters or less, but you cannot abbreviate the command prefix (SNAPMGR or NDMGR).

For example, these commands are all the same:

```
SnapMgr> SNAPSHOT "My Snapshot" /DELETE
SnapMgr> SNAPSHO "My Snapshot" /DEL
SnapMgr> SNA "My Snapshot" /D
```

5.2.5 Using Uppercase or Lowercase in Commands

You can type the commands in uppercase, lowercase, or a mixture of both. For example, these commands are all the same:

```
SnapMgr> DRIVES
SnapMgr> Drives
SnapMgr> drives
SnapMgr> drIVES
```

5.3 Security and Privileges

To use the snap-ins or commands, your user account, or a group of which you are a member, must be in the local Administrators group on the computer you want to manage. If you want to manage a cluster, it must be in the local Administrators group on **every node** in the cluster.

By default on Windows NT, the Domain Admins global group is a member of the local Administrators group on every computer. So by default, any user in the Domain Admins group can use the snap-ins and commands to manage any computer or cluster in the domain.

If you want to allow users to manage the SWVR software without giving them the broad privileges associated with the Domain Admins global group:

1. Create a global group called SCE Admins. A domain administrator must create this group.
2. Add the users to the SCE Admins global group.
3. Add the SCE Admins global group to the local Administrators group on each computer that you want the users to manage. Remember to do this on every node in a cluster.

5.3.1 Managing other domains

You can also manage a computer in another domain provided that:

- There is a trust relationship between the two domains, **and**
- Your user account is in the local Administrators group on every computer you want to manage, including every node in a cluster.

6 Managing Pools

This chapter describes how to create and manage pools, virtual disks, and snapshots. The following topics are described in this chapter:

- Creating Pools, Virtual Disks, and Snapshots (Page 6-2)
 - Pool Space (Page 6-9)
 - Understanding Delspace (Page 6-11)
 - Adding Storage Units to a Pool (Page 6-12)
 - Restoring Virtual Disks (Page 6-15)
 - Using Snapshots to do Backups (Page 6-16)
 - Defragmenting Virtual Disks (Page 6-18)
 - Mapping a Drive Letter to a Virtual Disk or Snapshot (Page 6-18)
 - Partitioning and Formatting a Virtual Disk (Page 6-19)
 - Unmapping a Drive Letter from a Virtual Disk or Snapshot (Page 6-20)
 - Displaying Information about Pools (Page 6-21)
 - Displaying Information about Virtual Disks and Snapshots (Page 6-25)
 - Deleting a Snapshot (Page 6-27)
 - Deleting a Virtual Disk (Page 6-28)
 - Deleting a Pool (Page 6-29)
- ▲ **Tip:** Whenever there are two ways to do a task – using Snapshot Manager or using another tool – always use Snapshot Manager. For example, use Snapshot Manager to map drive letters to virtual disks and snapshots, not Disk Administrator.

6.1 Creating Pools, Virtual Disks, and Snapshots

The following sections describe how to create pools, virtual disks, and snapshots.

6.1.1 Creating a Pool

Before you create a pool, decide which storage units you want to use. Note the following:

- You can have up to 8 storage units in a pool.
- They can be standard single-spindle disks or controller-based fault-tolerant disk sets. They can have different capacities and be from different manufacturers.
- They should all have the same read-write, redundancy and failure characteristics. For example, they should all be standard disks, or they should all be StorageWorks RAID 5 arrays, or they should all be mirror sets.
- In a cluster, use disks only on the shared storage bus; do not use local disks.
- Do not use network disks that are being served by remote computers or clusters.

You can have any number of pools on a standalone computer or cluster. The number is limited only by the number of storage units that are available.

- ▲ **Tip:** In a cluster, you **must** use Snapshot Manager to create pools, as described in this section. Snapshot Manager automatically creates cluster resources of type SCE Pool and SCE Storage Unit and brings them online. In a cluster, you cannot use Cluster Administrator to create resources of type SCE Pool or SCE Storage Unit.
- ▲ **Tip:** In a cluster, you cannot create a pool if any node in the cluster is down.

Steps

To create a pool:

1. If any of the storage units you want to use in the pool contains any data you want to keep, back up the data.
2. In a cluster, make sure that the storage units are not configured as cluster resources of type Physical Disk.

Use Cluster Administrator to check this, and to delete the resources if necessary.

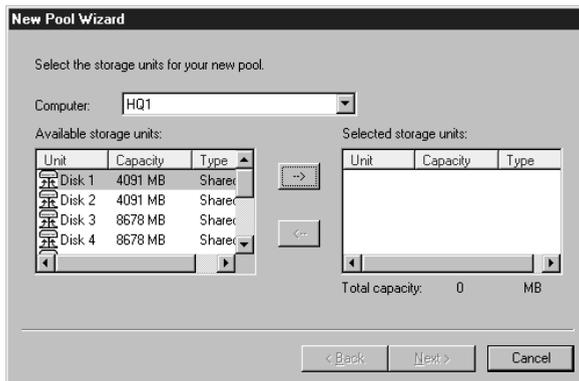
3. If any of the storage units contains a partition, use Disk Administrator to delete the partition.
4. Run the MMC Management Console.

Click **Start**, then **Programs**, then **Compaq StorageWorks Virtual Replicator**, and then **Management Console**.

5. In the left-hand Scope pane, select the **Snapshot Manager** snap-in for the standalone computer or cluster on which you want to create the pool.

If the snap-in has not yet been added to your console, add it now. For instructions on how to do this, see “Adding a Snap-In” on Page 5-4.

6. With the snap-in highlighted, right-click **New** then **Pool**.
7. You see the welcome screen of the New Pool wizard. Click **Next** to start the wizard.
8. The wizard displays a list of the available storage units.



Select the storage units you want to use, then click **Next**.

To select a storage unit, highlight it and then click the Right Arrow to move it to the right-hand side of the screen.

The wizard shows only non-removable disks that are online and have no partitions. It does not show virtual disks or snapshots; you cannot use them as storage units in pools.

In a cluster, the wizard shows the disks that can be seen by the node selected in the pull-down list box. By default, the selected node is the one that currently owns the cluster name resource.

The wizard shows the following information about each storage unit.

Unit	The disk number of the storage unit. In a cluster, this is the disk number on the node selected in the pull-down list box.
Capacity	The capacity of the storage unit.
Type	The type of the storage unit: Shared: a disk in a cluster that is on a shared storage bus and can be seen by every node in the cluster. Local: either a physical disk that is attached to the computer, or a network disk.
Total capacity	The combined capacities of all the selected storage units. The capacity of the pool will be slightly less than this, because some of the disk space is used to store internal pool configuration data. The capacity of the pool will be about 10% less than the total capacity displayed.

▲ **Tip:** In a cluster, select storage units that have only **Shared** in the **Type** column.

9. Type the name you want to give to the new pool.

The name can be up to 23 characters. Choose a name that is different from that of any other pool, virtual disk, or snapshot on the computer or cluster you are managing.

In a cluster, the name must also be different from that of any existing cluster group or cluster resource.

10. Choose a segment size, from 32K to 128K. The default is 32K. The segment size determines the largest possible virtual disk you can create in a pool. Then click **Next**.

11. Click **Finish** to complete the New Pool wizard.

The wizard creates the pool and marks the storage units as offline so that users can no longer access them directly.

In a cluster

The wizard automatically creates a cluster group that has the same name as the pool. Then it creates resources in the pool group:

- It creates a resource of type SCE Storage Unit for every storage unit in the pool. The name of the resource is Storage_{N}, where N is a 32-digit hexadecimal number that uniquely identifies the storage unit.

- It creates a resource of type SCE Pool for the pool itself. This resource has the same name as the pool. It is dependent on all of the storage unit resources. A pool must always be able to access all of its storage units, so it can't come online until all of the storage units are online.

The wizard brings the pool group and all of these resources online.

- ▲ **Note:** Do not use Cluster Administrator to rename any of these resources, or remove the dependency of the pool resource on the storage unit resources, because you may not be able to access your data.

6.1.2 Creating a Virtual Disk

Once you have created a pool, you can create virtual disks in it, of whatever capacities you want. The capacity of each virtual disk must be at least 10 MB. The segment size determines the largest possible virtual disk you can create in a pool. You can create up to 8 virtual disks in each pool.

Steps

To create a virtual disk in a pool:

1. Run the MMC Management Console.
Click **Start**, then **Programs**, then **Compaq StorageWorks Virtual Replicator**, and then **Management Console**.
2. In the left-hand Scope pane, select the **Snapshot Manager** snap-in for the cluster or standalone computer that the pool is on.
If the snap-in has not yet been added to your console, add it now. For instructions on how to do this, see "Adding a Snap-In" on Page 5-4.
3. Select the pool, and then right-click **New** then **Virtual Disk**.
4. You see the welcome screen of the New Virtual Disk wizard. Click **Next** to start the wizard.
5. Type the name you want to give the new virtual disk, and the capacity you want it to have, then click **Next**.

The name can be up to 23 characters. Choose a name that is different from that of any other pool, virtual disk, or snapshot on the computer or cluster you are managing.

In a cluster, if you want to serve the virtual disk as a network disk, choose a name that is different from that of any cluster resource, otherwise when you subsequently serve it, you will not be able to make its network disk name the same as its virtual disk name.

The capacity of the virtual disk must be at least 10 MB, and cannot exceed the free space in the pool, which is displayed by the wizard.

6. Optionally, map a drive letter to the new virtual disk.

To map a drive letter, select a letter from the pull-down list, then click **Next**. In a cluster, the list only shows drive letters that are available on every node in the cluster that is currently online.

If possible, select a letter near the end of the alphabet.

During system startup, drive letters are automatically allocated from the start of the alphabet. When the pool software starts up, no drive letter is allocated to the virtual disk if the drive letter you chose has already been allocated to another disk that came online earlier in the startup sequence. The pool software starts quite late in the startup sequence, so select a letter near the end of the alphabet to reduce the probability that the drive letter is allocated to another disk.

In a cluster, if the drive letter is not available when the pool fails over to another node, no drive letter is mapped to the virtual disk.

If you do not want to map a drive letter now, click the **No drive letter** button, then click **Next**. The virtual disk will be offline until you map a drive letter to it. When you subsequently map a drive letter, do not use Disk Administrator. Follow the instructions in “Mapping a Drive Letter to a Virtual Disk or Snapshot” on Page 6-18.

7. If you mapped a drive letter to the virtual disk in Step 6, you now have the option of formatting the disk.

Click **Next** to format the disk. This creates a single partition on the disk and formats it with the NTFS file system. It gives the partition the same volume label as the virtual disk name.

If you do not want to format the disk now, click **No**, and then click **Next**.

8. Click **Finish** to complete the wizard.

In a cluster

The wizard does not create any cluster resources. There are cluster resources for pools and their storage units, but not for virtual disks or snapshots. The virtual disks and snapshots in a pool automatically appear on whichever node currently owns the pool resource.

6.1.3 Creating a Snapshot

You can create a snapshot of either a virtual disk or another snapshot. The disk you create the snapshot of is called the **parent** disk. The snapshot is in the same pool, and in the same family within that pool, as its parent disk. You can have up to 12 snapshots in each family.

Creating a snapshot flushes the local system cache. Any data in the cache that has not yet been written out to the parent disk is flushed to disk before the snapshot is created.

- ▲ **Tip:** If the parent disk is being served as a network disk, it does not have any data in the local cache, but may have data in the cache on the remote computer. To flush the remote cache, disconnect the remote computer from the parent disk before you create the snapshot.

Steps

To create a snapshot of a disk in a pool:

1. Run the MMC Management Console.
Click **Start**, then **Programs**, then **Compaq StorageWorks Virtual Replicator**, and then **Management Console**.
2. In the left-hand Scope pane, select the **Snapshot Manager** snap-in for the standalone computer or cluster that the pool is on.
If the snap-in has not yet been added to your console, add it now. For instructions on how to do this, see “Adding a Snap-In” on Page 5-4.
3. Select the virtual disk or snapshot that you want to create a snapshot of, and then right-click **New** then **Snapshot**.
4. You see the welcome screen of the New Snapshot wizard. Click **Next** to start the wizard.
5. Type the name you want to give the new snapshot, then click **Next**.
The name can be up to 23 characters. Choose a name that is different from that of any other pool, virtual disk, or snapshot on the computer or cluster you are managing.
In a cluster, if you want to serve the snapshot as a network disk, choose a name that is different from that of any cluster resource, otherwise when you subsequently serve it, you will not be able to make its network disk name the same as its snapshot name.
6. Optionally, map a drive letter to the new snapshot.
To map a drive letter, click the **Yes** button, select a drive letter from the pull-down list, and then click **Next**. In a cluster, the list shows only drive letters that are available on every node in the cluster that is currently online. If possible, select a letter near the end of the alphabet.

During system startup, drive letters are automatically allocated from the start of the alphabet. When the pool software starts up, no drive letter is allocated to the snapshot if the drive letter you chose has already been allocated to another disk that came online earlier in the startup sequence. The pool software starts quite late in the startup sequence, so select a letter near the end of the alphabet to reduce the probability that the drive letter is allocated to another disk.

In a cluster, if the drive letter is not available when the pool fails over to another node, no drive letter is mapped to the snapshot.

If you do not want to map a drive letter now, accept the default (**No drive letter**), and just click **Next**. The snapshot will be offline until you map a drive letter to it. When you subsequently map a drive letter, do not use Disk Administrator. Follow the instructions in “Mapping a Drive Letter to a Virtual Disk or Snapshot” on Page 6-18.

7. Click **Finish** to complete the wizard.

- ▲ **Tip:** To help you recognize the snapshot disk in Windows NT Explorer, change its volume label to be the same as its snapshot name. By default, it has the same volume label as its parent because it is an identical copy of its parent, which can be confusing.

In a cluster

The wizard does not create any cluster resources. There are cluster resources for pools and their storage units, but not for virtual disks or snapshots. The virtual disks and snapshots in a pool automatically appear on whichever node currently owns the pool resource.

6.1.4 Scheduling Automatic Snapshots

The virtual disk property sheet contains a property page, **Snapshot Schedules**, on which you can specify:

- When to take periodic snapshots of the virtual disk.
- How many snapshots to retain over time.

When the total number of snapshots to be retained has been reached or exceeded, SWVR deletes the oldest periodic snapshot not mapped to a drive letter. If all snapshots are mapped to drive letters, SWVR creates a single snapshot. This is the only time the number of snapshots that should be maintained is exceeded. If the total number of snapshots for the virtual disk reaches the SWVR limit of 12, and no periodic snapshots can be deleted, then no further snapshots are taken.

6.2 Pool Space

The following sections describe how to monitor disk space in a pool.

6.2.1 Monitoring Pool Free Space

It is essential that your pools do not run out of disk space. If a pool runs out of disk space, writes to it fail and you may lose data. Use Snapshot Planner before you set up your pools and virtual disks. Snapshot Planner cannot see virtual disks, so once you have created virtual disks in a pool, you cannot use it to predict the cost of snapshots of your virtual disks.

When a pool runs out of disk space, you must either add a storage unit to it, or delete one of its snapshots or virtual disks.

- ▲ **Tip:** If a pool is full, do not try to free up space by deleting files. Deleting files does not free up any space in the pool.

If the capacity of one of your pools is less than the combined capacities of all of its virtual disks and snapshots, there is a possibility that it could become full. In this case, you should regularly monitor its free space.

- ▲ **Tip:** To monitor pool free space, you **must** use the Snapshot Manager snap-in or commands.

To monitor a pool's free space:

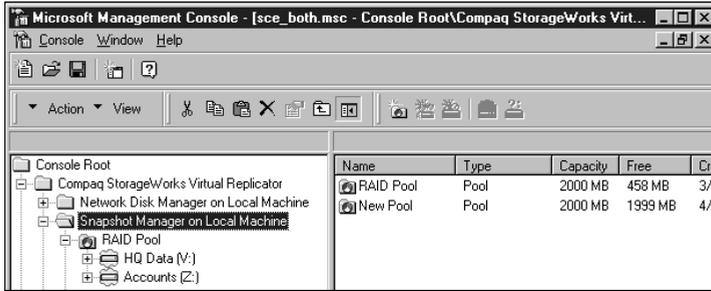
1. Run the MMC Management Console.

Click **Start**, then **Programs**, then **Compaq StorageWorks Virtual Replicator**, and then **Management Console**.

2. In the left-hand Scope pane, select the **Snapshot Manager** snap-in for the standalone computer or cluster that the pool is on.

If the snap-in has not yet been added to your console, add it now. For instructions on how to do this, see "Adding a Snap-In" on Page 5-4.

- When the snap-in is selected in the left-hand Scope pane, the right-hand Results pane shows all the pools that are currently online on the computer or cluster. It shows the capacity and free space of each pool.



In this example, the RAID pool has a capacity of 2000 MB but it has only 458 MB of free space. It is close to being full.

When the free space falls to less than 30%, the *Virtual Replicator* software starts logging warning records every 5 minutes in the system event log to warn you that the pool is getting full. Then it logs error records every 5 minutes that say that the pool's free space is less than 5%. These records have **sddriver** in the source column.

When you select the icons for the virtual disks and snapshots in the pool and right-click **Properties**, you see three pie charts, which are summarized in this table.

Disk	Capacity	Space used	Free space
VirtualDisk (Z:)	3,000 MB	2,100 MB	900 MB
VirtualDisk Copy 1 (Y:)	3,000 MB	2,250 MB	750 MB
VirtualDisk Copy 2 (X:)	3,000 MB	2,000 MB	1,000 MB
Totals	9,000 MB	6,350 MB	2,650 MB

According to Explorer, the combined free space on the three disks is 2,650 MB, when in fact there is only 4 MB of free space.

If you deleted files on any of the disks, for example, if you deleted files on drive Z: (VirtualDisk), you would see the free space rising in Windows NT Explorer, but not in Snapshot Manager. Deleting files does not free up space in the pool.

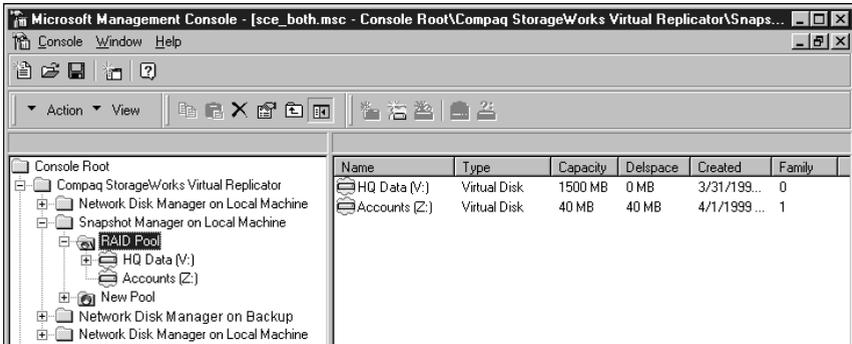
To increase the pool free space you **must** either:

- Add a storage unit to the pool, or
- Delete one of its snapshots or virtual disks

6.2.2 Understanding Delspace

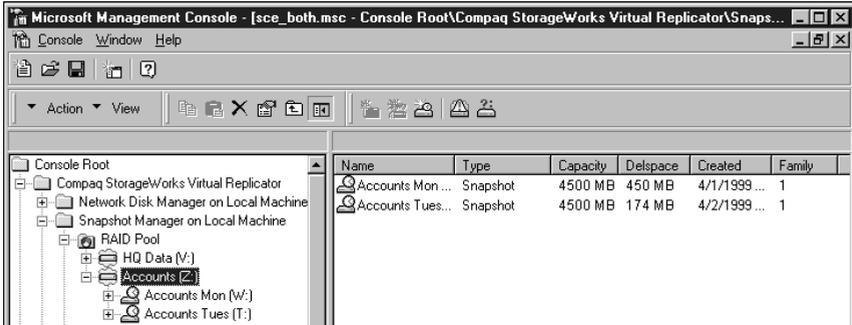
The **Delspace** of a virtual disk or snapshot is the amount of space you free in the pool if you delete it. It is the amount by which the free space in the pool increases if you delete that virtual disk or snapshot now.

To display the Delspace of all the virtual disks in a pool, select the pool in the left-hand Scope pane of the Snapshot Manager snap-in. The right-hand Results pane shows all the virtual disks in the pool.



The Delspace of every virtual disk is equal to either 0 or the capacity of the virtual disk. If a virtual disk has a snapshot, its Delspace is 0. That is because you cannot delete it if there are any snapshots in its family. As soon as you delete the last snapshot, the Delspace of the virtual disk changes from 0 to its capacity.

To display the Delspace of a snapshot, select the virtual disk for the snapshot's family in the left-hand Scope pane. The right-hand Results pane shows all the snapshots in the family.



In this example, the Delspace for the snapshot Accounts Mon is 450 MB. If you deleted Accounts Mon, the free space in the pool would increase by 450 MB.

The Delspace of a snapshot is the amount of space used exclusively by that snapshot.

When there is only one snapshot in a family, its Delspace is actually the amount of space that it is using in the pool.

But when there are two or more snapshots in a family, the snapshots can share disk space with each other, and so you cannot tell how much space each individual snapshot is using.

And when you delete two snapshots in the same family, you may free up more space than the sum of their Delspaces. For example, when you delete Accounts Mon and Accounts Tue you free up 1000 MB, although the sum of their Delspaces is 624 MB. The difference between the amount you free up and the sum of the Delspaces (376 MB) is the amount of disk space that Accounts Mon shared with Accounts Tue.

Suppose you deleted Accounts Mon first. As soon as you deleted Accounts Mon, Accounts Tue's Delspace would increase from 174 MB to 550 MB, because the 376 MB of shared disk space are now used exclusively by Accounts Tue.

6.3 Adding Storage Units to a Pool

You can add storage units to a pool at any time, even when users are accessing its virtual disks and snapshots.

Before you add the storage units, decide which units you want to use. Remember:

- You can have up to 8 storage units in a pool.
- They can be standard single-spindle disks or controller-based fault-tolerant disk sets. They can have different capacities and be from different manufacturers.
- They should all have the same read-write, redundancy and failure characteristics as the existing storage units in the pool. For example, if the existing storage units are StorageWorks RAID 5 arrays, you should only add StorageWorks RAID 5 arrays.
- In a cluster, only use disks on the shared storage bus; do not use local disks.
- Do not use network disks that are being served by remote computers or clusters.

Tip: In a cluster, you cannot add a storage unit to a pool if any of the nodes in the cluster are down.

Steps

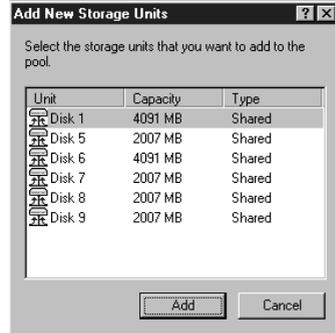
To add storage units to a pool:

1. If any of the storage units contains any data you want to keep, back up the data.
2. In a cluster, make sure that the storage units are not configured as cluster resources of type Physical Disk.
Use Cluster Administrator to check this, and to delete the resources if necessary.
3. If any of the storage units contains a partition, use Disk Administrator to delete the partition.
4. Run the MMC Management Console.
Click **Start**, then **Programs**, then **Compaq StorageWorks Virtual Replicator**, and then **Management Console**.
5. In the left-hand Scope pane, select the **Snapshot Manager** snap-in for the standalone computer or cluster on which the pool resides.
If the snap-in has not yet been added to your console, add it now. For instructions on how to do this, see “Adding a Snap-In” on Page 5-4.
6. Select the pool, then right-click **Properties** to display the Pool Properties screen.

- Click the **Storage Units** tab, and then click **Add** to display the Add New Storage Units screen.

The screen shows only non-removable disks that are online and have no partitions. It does not show virtual disks or snapshots; you cannot use them as storage units in pools.

In a cluster, the screen shows the disks that can be seen by the node that currently owns the pool resource.



It shows the following information about each storage unit.

Unit	The disk number of the storage unit. In a cluster, this is the disk number on the node that currently owns the pool resource.
Capacity	The capacity of the storage unit.
Type	The type of the storage unit: Shared: a disk in a cluster that is on a shared storage bus and can be seen by every node in the cluster. Local: either a physical disk that is attached to the computer, or a network disk.

- Select the storage units you want to add, then click **Add**.
 - ▲ **Tip:** In a cluster, select only storage units that have **Shared** in the **Type** column.
- Click **OK** to dismiss the popup that tells you that the storage units have been added to the pool.
- Click **OK** to dismiss the Pool Properties screen.

The storage units are automatically marked as offline so that users can no longer access them directly.

In a cluster

For each storage unit that you added, Snapshot Manager automatically:

- Creates a resource of type SCE Storage Unit in the cluster group that contains the pool resource. The name of the new resource is `Storage_{N}`, where N is a 32-digit hexadecimal number that uniquely identifies the storage unit.
- Changes the properties of the pool resource to make it dependent on all of the new storage units resources.
 - ▲ **Note:** Do not use Cluster Administrator to rename the storage unit resources, or remove the dependency of the pool resource on them, because you may not be able to access your data.

6.4 Restoring Virtual Disks

In the event that the data on a virtual disk has become corrupted, there are three methods by which you can perform an online restore of the data:

- **Method 1:** Unmap the drive letter of the virtual disk and map the same drive letter to the snapshot.
- **Method 2:** Assign a drive letter to the snapshot, and copy files from the snapshot to the virtual disk. Take a new snapshot of the virtual disk, then delete the old snapshot.
- **Method 3:** Create a new virtual disk, and copy files from the snapshot to the new virtual disk. Once the copy is completed, delete the old virtual disk and all the snapshots.

Although Method 1 is the quickest way to make data available to users, it leaves the original virtual disk as an orphan, and causes some performance issues because of the copy-outs.

Recommendations:

- Use Method 2 to restore individual files that have been deleted or corrupted.
- Use Method 3 to restore an entire volume.

Note: Method 3 does cause temporary space consumption, because of the two virtual disks and the need to preallocate pool space.

If you have *Compaq Batch Scheduler* installed, you can schedule the Method 3 type of online restore process using the Restore For Backup wizard. See Chapter 8 for a description of this wizard.

6.5 Using Snapshots to do Backups

The following steps show how to use the SnapMgr and NDMgr commands, so that you can use batch jobs to automate your backups. You can use either the commands or the snap-ins.

Steps

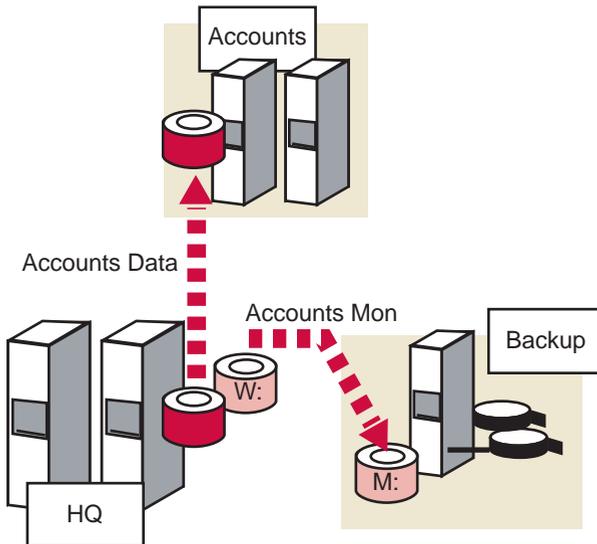
Follow these steps to use a snapshot to do a backup:

1. If you want the backup to capture the recent updates that applications are buffering in their own in-memory caches, shut down or quiesce the applications that are using the disk you want to back up.
2. If the disk you want to back up is being served as a network disk, and you want the backup to capture updates that are being buffered in the system cache on the remote client computer, disconnect the client.

For example, if you want to back up the virtual disk Accounts Data, which the HQ cluster is serving to the Accounts cluster:

```
C:\> NDMGR Accounts USE \\HQ\“Accounts Data” /DELETE
```

The following figure illustrates the example used in this section.



3. Create a snapshot of the disk you want to back up.

For example, this command creates a snapshot called Accounts Mon of the virtual disk Accounts Data on the HQ cluster:

```
C:\> SNAPMGR HQ SNAPSHOT “Accounts Mon” /PARENT:“Accounts Data”
```

Creating the snapshot flushes the local system cache. Any data in the cache that has not yet been written out to the parent disk is flushed to disk before the snapshot is created.

4. If you disconnected the remote client in Step 2, reconnect it now.

For example:

```
C:\> NDMGR Accounts USE * \\HQ\Accounts Data"
```

Use an asterisk as the first parameter of the USE command.

5. If you shut down or quiesced any applications in Step 1, restart them now.
6. Map a drive letter to the snapshot.

For example, map drive W to the snapshot:

```
C:\> SNAPMGR HQ SNAPSHOT "Accounts Mon" /MAP:W
```

7. If you want to do the backup locally, go to Step 10.

If you want to do the backup remotely, go to Step 8.

8. Serve the snapshot as a network disk. Make the network disk name the same as the name of the snapshot, for example:

```
C:\> NDMGR HQ SERVE "Accounts Mon" W
```

9. Connect the remote computer or cluster to the served snapshot.

For example, connect the Backup computer to the served snapshot and map drive letter M to it:

```
C:\> NDMGR Backup USE M: \\HQ\Accounts Mon"
```

You must issue this command on the Backup computer. If you issue it remotely, you cannot control which drive letter is mapped to the snapshot – the first available drive letter is mapped instead of the one you specify on the USE command.

- ▲ **Tip:** If you are running the backup in a batch job, run the batch job on the computer that does the backup. In this example, run the batch job on Backup.

10. Back up the snapshot.

For example, if the drive letter mapped to the snapshot is M:

```
C:\> NTBACKUP BACKUP M: /T NORMAL /TAPE:0
```

If your backup tool records information on the disk being backed up, remember to do a full backup, and not an incremental or differential backup (see “Existing Backup Tools?” on Page 3-7).

11. If you are doing the backup locally, go to Step 14.

If you are doing the backup remotely, go to Step 12.

12. Disconnect the remote computer or cluster from the served snapshot.

For example:

```
C:\> NDMGR Backup USE \\HQ\“Accounts Mon” /DELETE
```

13. Stop serving the snapshot.

For example:

```
C:\> NDMGR HQ SERVE “Accounts Mon” /DELETE
```

14. Unmap the drive letter from the snapshot then delete the snapshot.

For example:

```
C:\> SNAPMGR HQ SNAPSHOT “Accounts Mon” /UNMAP /DELETE
```

6.6 Defragmenting Virtual Disks

Before you defragment a virtual disk, you delete all the snapshots in its family to avoid unnecessary copy-out operations. This improves performance and avoids using up disk space.

Defragmenting a disk consolidates the space used by the files on the disk. This involves moving data about on the disk.

If you are defragmenting a virtual disk that has a snapshot, the process of moving the data causes copy-out operations. Although the data is not **changing**, the contents of sectors on the disk are changing because the data is being **moved** to new locations on disk. This causes unnecessary copy-out operations that degrade performance and waste disk space.

6.7 Mapping and Unmapping Drives

The following sections describe how to map a drive letter, partition and format a virtual disk, and unmap a drive letter.

6.7.1 Mapping a Drive Letter to a Virtual Disk or Snapshot

Use Snapshot Manager, not Disk Administrator, to map drive letters to virtual disks and snapshots.

Steps

To map a drive letter to a virtual disk or snapshot:

1. Run the MMC Management Console.
Click **Start**, then **Programs**, then **Compaq StorageWorks Virtual Replicator**, and then **Management Console**.
2. In the left-hand Scope pane, select the **Snapshot Manager** snap-in for the standalone computer or cluster that the virtual disk or snapshot is on.
If the snap-in has not yet been added to your console, add it now. For instructions on how to do this, see “Adding a Snap-In” on Page 5-4.
3. Select the virtual disk or snapshot, and then right-click **Task** then **Map/Unmap Drive Letter**.
4. Select a drive letter from the pull-down list, then click **Next**.
In a cluster, the list shows drive letters that are available only on every node in the cluster that is currently online.
If possible, select a letter near the end of the alphabet.
During system startup, drive letters are automatically allocated from the start of the alphabet. When the pool software starts up, if the drive letter you selected has already been allocated to another disk that came online earlier in the startup sequence, no drive letter is allocated to the virtual disk or snapshot. The pool software starts quite late in the startup sequence, so select a letter near the end of the alphabet to reduce the probability that the drive letter is allocated to another disk.
In a cluster, if the drive letter is not available when the pool fails over to another node, no drive letter is mapped to the virtual disk or snapshot.
5. On the next screen, click **Finish**.
The drive letter is mapped to the virtual disk or snapshot, and then the disk is brought online.

6.7.2 Partitioning and Formatting a Virtual Disk

You can use either Disk Administrator or Snapshot Manager to partition and format a virtual disk. The following steps describe how to use Snapshot Manager.

If you use Disk Administrator, make sure that you select the NTFS file system. *Compaq StorageWorks Virtual Replicator* does not support the FAT file system.

- ▲ **Tip:** You must map a drive letter to the virtual disk before you can format it. If it does not have a drive letter, follow the instructions in “Mapping a Drive Letter to a Virtual Disk or Snapshot” on Page 6-18.

Steps

To partition and format a virtual disk:

1. Run the MMC Management Console.
Click **Start**, then **Programs**, then *Compaq StorageWorks Virtual Replicator*, and then **Management Console**.
2. In the left-hand Scope pane, select the **Snapshot Manager** snap-in for the standalone computer or cluster that the virtual disk is on.
If the snap-in has not yet been added to your console, add it now. For instructions on how to do this, see “Adding a Snap-In” on Page 5-4.
3. Select the virtual disk, and then right-click **Task** then **Format**.
4. Click **OK** to confirm that you want to format the disk and delete all the data on it.

This creates a single partition on the virtual disk and formats it with the NTFS file system. It gives the partition the same volume label as the virtual disk name.

6.7.3 Unmapping a Drive Letter from a Virtual Disk or Snapshot

Use Snapshot Manager, not Disk Administrator, to unmap drive letters from virtual disks and snapshots.

Steps

To unmap a drive letter from a virtual disk or snapshot:

1. Make sure that no-one is accessing the virtual disk or snapshot. You cannot unmap its drive letter if any files on it are open.
2. Run the MMC Management Console.
Click **Start**, then **Programs**, then *Compaq StorageWorks Virtual Replicator*, and then **Management Console**.

3. In the left-hand Scope pane, select the **Snapshot Manager** snap-in for the standalone computer or cluster that the virtual disk or snapshot is on.

If the snap-in has not yet been added to your console, add it now. For instructions on how to do this, see “Adding a Snap-In” on Page 5-4.

4. Select the virtual disk or snapshot, and then right-click **Task** then **Map/Unmap Drive Letter**.
5. Check that the **No drive letter** button is selected, then click **Next**.
6. On the next screen, click **Finish**.

The drive letter is unmapped from the virtual disk or snapshot, and then the disk is taken offline.

6.8 Displaying Information

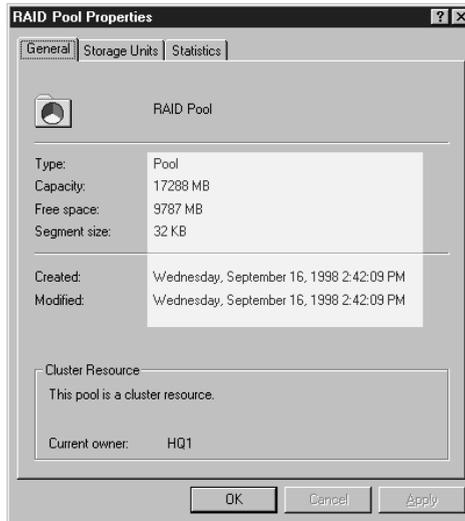
The following sections describe how to display information about pools, virtual disks, and snapshots.

6.8.1 Displaying Information about Pools

To display summary information about all the pools on a standalone computer or cluster, select the Snapshot Manager snap-in for the computer or cluster. The right-hand Results pane shows the following information about each pool; in a cluster, it does not show pools that are currently offline.

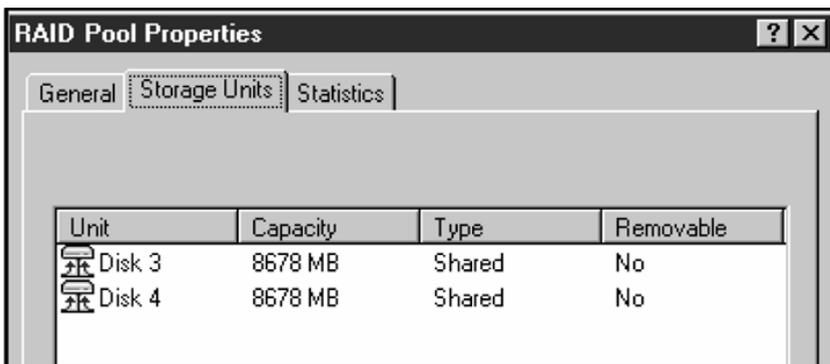
Column	Description
Name	The name of the pool.
Type	Pool.
Capacity	The capacity of the pool, rounded up to the nearest megabyte.
Free	The amount of free space currently available in the pool, rounded down to the nearest megabyte.
Created	When the pool was created.
Owner	The name of the cluster node that currently owns the pool resource. You do not see this column on a standalone computer.

For more information, select the pool then right-click **Properties**.



The **General** tab shows the following additional information:

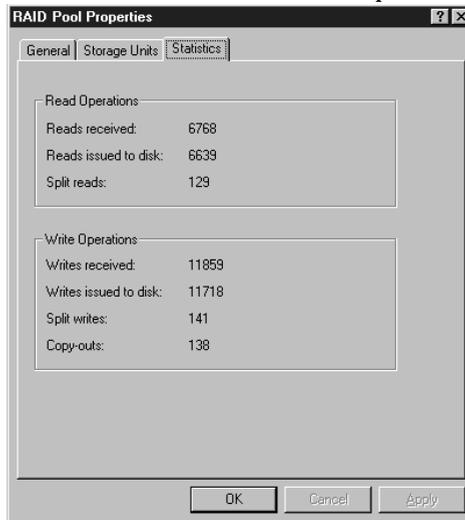
- The pool's segment size. A segment is the smallest unit that is copied during a copy-out operation.
- When it was last modified. This is when a storage unit was last added to it.



The **Storage Units** tab shows which storage units are in the pool. For each unit it shows the following information:

- Its disk number. In a cluster, this is the disk number on the node that currently owns the pool resource. The disk number may change when you restart the computer or when the pool resource fails over within the cluster.
- Its capacity.
- Its type:
 - Shared:** a disk in a cluster that is on a shared storage bus and can be seen by every node in the cluster.
 - Local:** a physical disk that is attached to the computer.
- Whether it can be removed from the pool.

The **Statistics** tab shows I/O statistics for the pool.



Managing Pools

It shows the following counters, which are zeroed when the computer restarts, or, in a cluster, when the pool fails over within the cluster.

Counter	Description
Reads received	The number of read I/O requests received by the pool software.
Reads issued to disk	The number of read I/Os issued to disk by the pool software. This includes reads caused by copy-out operations and split reads.
Split reads	The number of read I/O requests that had to be split into two or more I/Os to disk because the read crossed a segment boundary, and the next segment was not contiguous with the current segment.
Writes received	The number of write I/O requests received by the pool software.
Writes issued to disk	The number of write I/Os issued to disk by the pool software. This includes writes caused by copy-out operations and split writes.
Split writes	The number of write I/O requests that had to be split into two or more I/Os to disk because the write crossed a segment boundary, and the next segment was not contiguous with the current segment.
Copy-outs	The number of segments that were copied to preserve data for snapshots.

6.8.2 Displaying Information about Virtual Disks and Snapshots

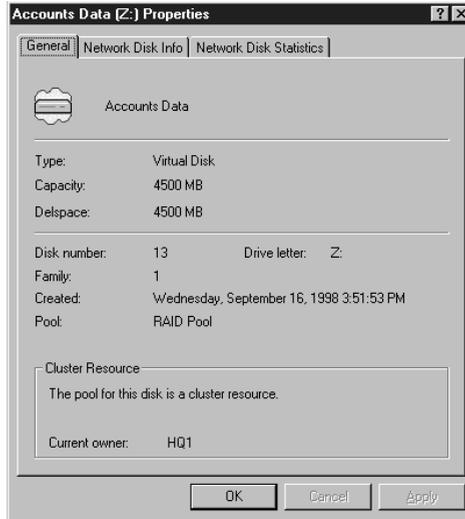
To display information about a virtual disk, select its pool in the left-hand Scope pane of MMC. The right-hand Results pane shows the following information about each virtual disk in the pool.

Column	Description
Name	The name of the virtual disk, and in brackets, the drive letter that is mapped to it.
Type	Virtual Disk.
Capacity	The capacity of the virtual disk.
Delspace	The amount of space you would free in the pool if you deleted the virtual disk. It is the same as the capacity of the virtual disk, unless there are any snapshots in its family, in which case it is 0. When a virtual disk has a snapshot in its family, you cannot delete the virtual disk, and so its Delspace is 0.
Created	When the virtual disk was created.
Family	In which family the virtual disk is.
Owner	The name of the cluster node that currently owns the pool resource. You do not see this column on a standalone computer.

To display information about a snapshot, select its family's virtual disk in the left-hand Scope pane of MMC. The right-hand Results pane shows the following information about each snapshot in the family.

Column	Description
Name	The name of the snapshot, and in brackets, the drive letter that is mapped to it.
Type	Snapshot.
Capacity	The capacity of the snapshot. This is how big the file system and every other piece of software on the system thinks the snapshot is. For example, it is the size shown by Windows NT Explorer. It is the maximum pool space the snapshot could consume, if you modified the entire contents of either the snapshot or its parent disk.
Delspace	The amount of space you would free in the pool if you deleted the snapshot right now. It is the amount of pool space used exclusively by this snapshot – the space that it is using and that is not shared with any other snapshots in its family. When there is only one snapshot in a family, the Delspace is the amount of pool space used by the snapshot.
Created	When the snapshot was created.
Family	In which family the snapshot is.
Owner	The name of the cluster node that currently owns the pool resource. You do not see this column on a standalone computer.

For more information, select the virtual disk or snapshot, then right-click **Properties**.



The **General** tab shows the disk number of the virtual disk or snapshot. The disk number may change when you restart the computer or when the pool resource fails over within the cluster.

If the virtual disk or snapshot is being served, there are **Network Disk Info** and **Network Disk Statistics** tabs, which show information such as the network disk name and the name of the client computer that is connected to it. For information on what these tabs display, see "Showing Information about Served Network Disks" on Page 7-17.

6.9 Deleting Snapshots, Virtual Disks, and Pools

The following sections describe how to delete snapshots, virtual disks, and pools.

6.9.1 Deleting a Snapshot

If you no longer need a snapshot, delete it to free space in the pool. The pool's free space increases by the snapshot's Delspace.

When you delete the snapshot, the Delspace of other disks in its family may rise:

- If it shared disk space with another snapshot, the Delspace of the other snapshot rises by the amount of the shared disk space.

- If it was the only snapshot in its family, the Delspace of the family's virtual disk changes from 0 to the capacity of the virtual disk. Now you can delete the virtual disk, so its Delspace becomes the same as its capacity.

Deleting the snapshot destroys all the data stored on the snapshot, but it does not affect the data stored on its parent disk or on any other disks in its family.

- ▲ **Tip:** You cannot delete a snapshot if you are serving it as a network disk.

Steps

To delete a snapshot:

1. If the snapshot contains any data you want to keep, back up the data. Deleting a snapshot destroys all the data stored on it.
2. Make sure that no-one is accessing the snapshot. You cannot delete it if any files on it are open.
3. Run the MMC Management Console.
Click **Start**, then **Programs**, then **Compaq StorageWorks Virtual Replicator**, and then **Management Console**.
4. In the left-hand Scope pane, select the **Snapshot Manager** snap-in for the standalone computer or cluster on which the snapshot is.
If the snap-in has not yet been added to your console, add it now. For instructions on how to do this, see “Adding a Snap-In” on Page 5-4.
5. Select the snapshot, then right-click **Delete**.
6. Click **OK** to confirm that you want to delete the snapshot and erase all the data stored on it.

6.9.2 Deleting a Virtual Disk

If you no longer need a virtual disk, delete it to free the space it is using in the pool.

Tip: Delete a virtual disk only after you have deleted all the snapshots in its family. Do not delete a virtual disk if you are serving it as a network disk.

Steps

To delete a virtual disk:

1. If the virtual disk contains any data you want to keep, back up the data. Deleting a virtual disks destroys all the data stored on the disk.

2. Make sure that no-one is accessing the virtual disk. You cannot delete it if any files on it are open.
3. Run the MMC Management Console.
Click **Start**, then **Programs**, then **Compaq StorageWorks Virtual Replicator**, and then **Management Console**.
4. In the left-hand Scope pane, select the **Snapshot Manager** snap-in for the standalone computer or cluster that the virtual disk is on.
If the snap-in has not yet been added to your console, add it now. For instructions on how to do this, see “Adding a Snap-In” on Page 5-4.
5. Select the virtual disk, then right-click **Delete**.
6. Click **OK** to confirm that you want to delete the virtual disk and erase all the data stored on it.

6.9.3 Deleting a Pool

If you no longer need a pool, delete it so that you can use its storage units for something else.

Deleting the pool frees its storage units and brings them online, so that you can access them directly. You can use them to create a new pool, or you can partition and format them, and use them as ordinary disks.

- ▲ **Tip:** You cannot delete a pool if it contains any virtual disks or snapshots.

Steps

To delete a pool:

1. Run the MMC Management Console.
Click **Start**, then **Programs**, then **Compaq StorageWorks Virtual Replicator**, and then **Management Console**.
2. In the left-hand Scope pane, select the **Snapshot Manager** snap-in for the standalone computer or cluster that the pool is on.
If the snap-in has not yet been added to your console, add it now. For instructions on how to do this, see “Adding a Snap-In” on Page 5-4.
3. Select the pool, then right-click **Delete**.
4. Click **OK** to confirm that you want to delete the pool.

In a cluster

Deleting the pool automatically deletes the pool resource (of type SCE Pool), and the storage unit resource (of type SCE Storage Unit) for each storage unit in the pool. It also deletes the pool group, if the group is now empty.

7 Managing Network Disks

This chapter describes how to manage network disks. The following topics are described in this chapter:

- Serving a Disk (Page 7-1)
- Controlling Which Clients Connect to the Disks (Page 7-7)
- Connecting to a Served Network Disk (Page 7-8)
- Changing the IP Subnet Configuration (Page 7-13)
- Changing the Client System for a Served Network Disk (Page 7-14)
- Using a Served Network Disk as the Cluster Quorum Disk (Page 7-14)
- Tuning Tips (Page 7-16)
- Showing Information about Served Network Disks (Page 7-17)
- Showing Information about Connected Network Disks (Page 7-19)
- Disconnecting from a Served Network Disk (Page 7-21)
- Stopping the Serving of a Network Disk (Page 7-22)
 - ▲ **Tip:** Whenever there are two ways to do a task – using Network Disk Manager or using Cluster Administrator – always use Network Disk Manager. For example, use Network Disk Manager, not Cluster Administrator, to create and delete cluster resources of type SCE Served Disk and SCE Connected Disk.

7.1 Serving a Disk

When you serve a disk as a network disk, the disk is marked as offline on the computer or cluster serving it. This prevents unsynchronized updates to the disk.

Users on the computer or cluster serving the disk can no longer access it directly. For example, if they double-click its icon in Windows NT

Explorer, they see a popup that says that the device is not ready. The disk remains offline until you stop serving it.

The computer or cluster automatically serves the disk whenever it restarts.

The steps you need to follow to serve a disk depend on which type of disk you are serving:

- **Serving Virtual Disks and Snapshots (Page 7-2)**
- **Serving other Disks (Page 7-3)**
 - ▲ **Tip:** In a cluster, do not use Cluster Administrator to create cluster resources of type SCE Served Disk. Use Network Disk Manager, as described in the following sections.

7.1.1 Serving Virtual Disks and Snapshots

To serve a virtual disk or snapshot as a network disk:

1. Make sure that all the files on the disk are closed. The disk is taken offline when you serve it.
2. Run the MMC Management Console.
Click **Start**, then **Programs**, then **Compaq StorageWorks Virtual Replicator**, and then **Management Console**.
3. In the left-hand Scope pane, select the **Snapshot Manager** snap-in for the standalone computer or cluster that the virtual disk or snapshot is on.
If the snap-in has not yet been added to your console, add it now. For instructions on how to do this, see “Adding a Snap-In” on Page 5-4.
 - ▲ **Tip:** You can also use the Network Disk Manager snap-in to serve a virtual disk or snapshot. It is quicker if you use the Snapshot Manager.
4. Select the virtual disk or snapshot you want to serve, and then right-click **Task** then **Serve as Network Disk**.
5. You see the welcome screen of the New Network Disk wizard. Click **Next** to start the wizard.
6. Click **Next** to accept the default name for the network disk. This is the same as the name of the virtual disk or snapshot you are serving.

7. On the next screen, you can change the list of registered clients.

This list affects every network disk served by the computer or cluster, not just the new disk you are serving. Any client on the list can connect to any network disk that is served by the computer or cluster. You cannot specify which clients can connect to which network disks.

Use **Add** and **Remove** to modify the list, then click **Next**.

- ▲ **Tip:** Use computer names, not cluster names, because cluster names are ignored in the list of registered clients. If you want to allow a cluster to connect to the network disk, add the name of every node in the cluster.

8. Click **Finish** to complete the wizard.

In a cluster

The wizard automatically creates a resource of type SCE Served Disk in the pool's cluster group. The resource has the same name as the network disk and is dependent on the pool resource.

The wizard brings the new resource online.

- ▲ **Note:** Do not use Cluster Administrator to change the dependency of the SCE Served Disk resource on the pool resource. If you use Cluster Administrator for this purpose, you may not be able to access your data.

7.1.2 Serving other Disks

To serve a disk other than a virtual disk or snapshot in a pool:

1. Make sure that all the files on the disk are closed. The disk is marked as offline when you serve it.
2. If the disk you want to serve is on a standalone computer, go to Step 4.
3. Make sure that the disk is not configured as a cluster resource of type Physical Disk.

Use Cluster Administrator to check whether the disk you want to serve is configured as a cluster resource of type Physical Disk and delete the resource if necessary.

4. Use Disk Administrator to map a drive letter to the disk and make it persistent. In a cluster, if the disk is on a shared storage bus, make sure that it has the same drive letter throughout the cluster.

Steps 5 to 7 describe how to do this.

5. Choose which drive letter you want to map to the disk.

In a cluster, use the NDMgr or SnapMgr DRIVES command (page B-2) to find out which drive letters are available throughout the cluster. Choose one of those drive letters.

6. Use Disk Administrator to map the drive letter to the disk and make it persistent.

Select the disk and deassign its current drive letter, even if it already has the same drive letter as the one you chose in Step 5. Now assign the drive letter you chose in Step 5 to the disk. The drive letter is now persistent on the local computer.

7. In a cluster, repeat Step 6 on every node that is attached to the shared storage bus. The disk now has a persistent drive letter throughout the cluster.

8. Run the MMC Management Console.

Click **Start**, then **Programs**, then **Compaq StorageWorks Virtual Replicator**, and then **Management Console**.

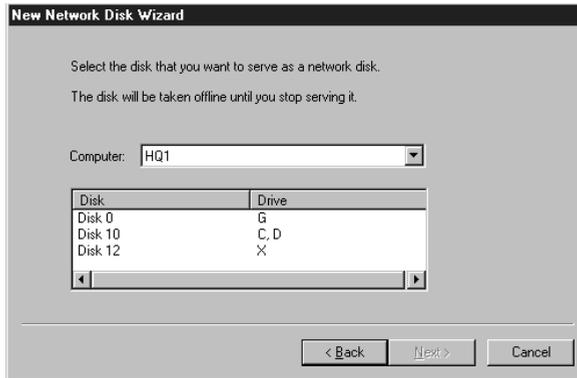
9. In the left-hand Scope pane, select the **Network Disk Manager** snap-in for the standalone computer or cluster on which the disk is.

If the snap-in has not yet been added to your console, add it now. For instructions on how to do this, see “Adding a Snap-In” on Page 5-4.

10. In the left-hand Scope pane, select the **Served Network Disks** folder, and then right-click **New** then **Network Disk**.

11. You see the welcome screen of the New Network Disk wizard. Click **Next** to start the wizard.

12. The wizard displays a list of disks. Select the one you want to serve, then click **Next**.



- ▲ **Tip:** The wizard displays all the hard disk drives on the system. This includes disks that you cannot serve, such as your system disk (in the previous screen capture, this is disk 10 and has two partitions, C and D), and disks that are already being served. If you select a disk you cannot serve, the final wizard screen will report an error saying it cannot serve that disk.
- ▲ **Tip:** Select a disk that has just one drive letter in the **Drive** column. *Compaq StorageWorks Virtual Replicator* does not support multiple partitions.

In a cluster, the wizard shows the disks that can be seen by the node selected in the pull-down list box. By default, the selected node is the one that currently owns the cluster name resource. Virtual disks and snapshots are treated like local disks – they are shown only when you select the node that currently owns the pool they are in.

13. Type the name you want to give to the network disk, then click **Next**.

This is the name the client uses when it connects to the disk. It can be up to 23 characters and is automatically converted to uppercase. Choose a name that is different from that of any other network disk that is being served by the computer or cluster.

In a cluster, the name must also be different from that of any existing cluster group or cluster resource.

- ▲ **Tip:** Make the name the same as the disk's volume label, so that you can recognize the disk when you use Windows NT Explorer on the client computer.

14. On the next screen, you can change the list of registered clients.

This list affects every network disk that is served by the computer or cluster, not just the new disk that you are serving. Any client on the list can connect to any network disk that is served by the computer or cluster. You cannot specify which clients can connect to which network disks.

Use **Add** and **Remove** to modify the list, then click **Next**.

- ▲ **Tip:** Use computer names, not cluster names, because cluster names are ignored in the list of registered clients. If you want to allow a cluster to connect to the network disk, add the name of every node in the cluster.

15. Click **Finish** to complete the wizard.

16. In a cluster, if you are serving a local disk, use the Cluster Administrator to restrict its possible owners to the system that is serving the local disk. For example, restrict the possible owners to the one computer that is attached to the disk.

- ▲ **Note:** If a failover is allowed to occur with a local disk and the failover system has a local disk with the same drive letter assigned, the wrong disk may be served.

In a cluster

The wizard automatically creates a cluster group and a cluster resource in the new group. Both the group and the resource have the same name as the network disk. The new resource is of type SCE Served Disk.

The wizard brings the new served disk resource online.

7.2 Controlling Which Clients Connect to the Disks

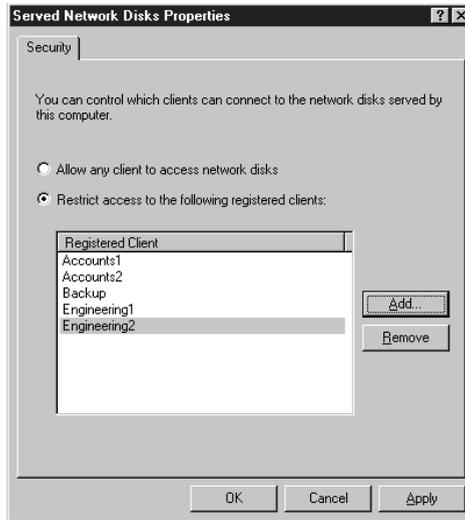
Each computer or cluster that serves network disks has a list of the client computers that are registered to connect to its disks.

Any computer on the list can connect to any of the network disks served by the computer or cluster. You cannot specify which computers can connect to which network disks.

Steps

To change the list of registered clients for a computer or cluster that serves network disks:

1. Run the MMC Management Console.
Click **Start**, then **Programs**, then **Compaq StorageWorks Virtual Replicator**, and then **Management Console**.
2. In the left-hand Scope pane, select the **Network Disk Manager** snap-in for the standalone computer or cluster whose list you want to change.
If the snap-in has not yet been added to your console, add it now. For instructions on how to do this, see “Adding a Snap-In” on Page 5-4.
3. In the left-hand Scope pane, select the **Served Network Disks** folder, then right-click **Properties**.
4. The Security screen is displayed.



Select the **Allow any** button if you want any computer to be able to connect to the network disks that are served by the computer or cluster you are managing.

To control access to the network disks that are served by the computer or cluster you are managing, make sure the **Restrict access** button is selected.

To remove a computer from the list, select the computer, then click **Remove**.

To add a computer to the list, click **Add**, then either type the name of a client computer or click **Browse** to browse for it on the network. If you type the name, the software doesn't check that the computer exists.

- ▲ **Tip:** Use computer names, not cluster names, because cluster names are ignored in the list of registered clients. If you want to allow a cluster to connect to the network disk, add the name of every node in the cluster.

5. Click **OK**. Then click **OK** again to save the changes you made.

7.3 Connecting to a Served Network Disk

Connecting a standalone computer or cluster to a network disk creates a new disk on the computer or cluster, which behaves like a disk that is physically attached to that computer or cluster. Applications running on it are not aware that the disk is actually located on the network.

The network disk is automatically reconnected whenever the computer or cluster restarts, until disconnect it, using either the USE command (Page B-23) or the snap-ins.

Only one computer at a time can connect to a network disk.

In a cluster, the network disk is automatically associated with a cluster resource of type SCE Connected Disk. Like all cluster resources, this resource can be online on only one node at any time. If that node fails, the resource automatically fails over to another node in the cluster.

- ▲ **Tip:** In a cluster, do not use Cluster Administrator to create resources of type SCE Connected Disk. Use Network Disk Manager, as described in this section.

Steps

To connect to a served network disk:

1. Run the MMC Management Console.

Click **Start**, then **Programs**, then **Compaq StorageWorks Virtual Replicator**, and then **Management Console**.

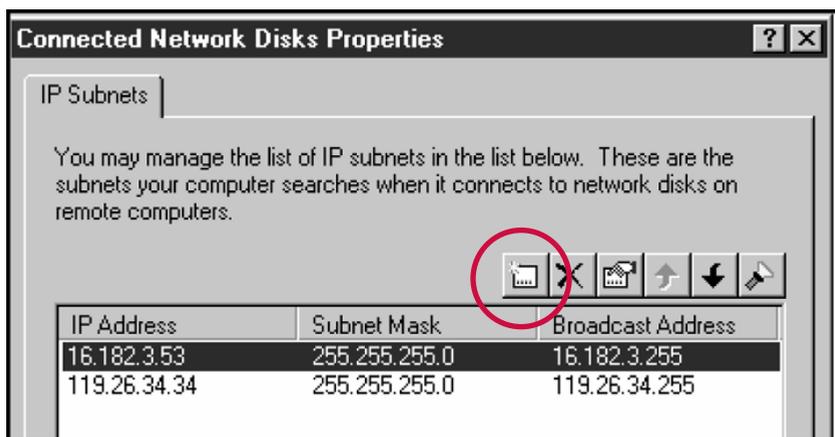
2. In the left-hand Scope pane, select the **Network Disk Manager** snap-in for the standalone computer or cluster that you want to connect to the network disk.

If the snap-in has not yet been added to your console, add it now. For instructions on how to do this, see “Adding a Snap-In” on Page 5-4.

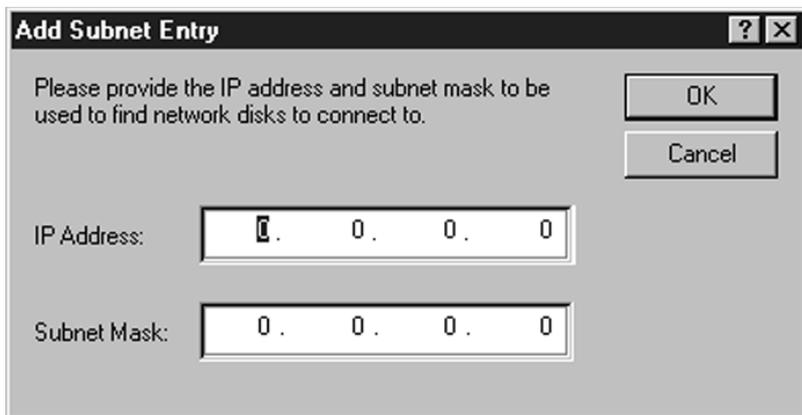
3. Go to Step 7 if the computer or cluster that you want to connect to the network disk is either in the same IP subnet as the one serving the network disk, or if it is already configured to connect to network disks in that subnet.

Otherwise follow Steps 4 to 6 to configure it to connect to network disks in that IP subnet.

- In the left-hand Scope pane, select the **Connected Network Disks** folder, then right-click **Properties**. You see the IP Subnets screen.



- Click the **New Subnet** toolbar button. You see the Add Subnet Entry screen.



6. Are your gateways configured to broadcast messages to other IP subnets?

Yes	Specify the subnet of the computer or cluster that is serving the disk by entering an IP address and subnet mask, then click OK twice.
No	Enter the IP address of the computer serving the disk, and a subnet mask of 255.255.255.255, then click OK . If the disk is being served by a cluster, repeat Steps 5 and 6 for each computer that can serve it, then click OK .

7. In the left-hand Scope pane, select the **Connected Network Disks** folder, and then right-click **New** then **Connected Network Disk**.

8. You see the welcome screen of the Network Disk Connection wizard. Click **Next** to start the wizard.

9. Provide the name of the computer or cluster that is serving the disk, then click **Next**.

You can either type the name or browse for it on the network. If the disk is being served by a cluster, you can type or select either the name of the cluster or the name of any computer in the cluster – the result is the same.

10. The wizard displays a list of all the disks that are being served by the computer or cluster that you selected in the previous step. It tells you which are available and which are not available.

Select the network disk to which you want to connect, then click **Next**.

11. Click **Finish** to complete the wizard.

12. The wizard automatically mapped a temporary drive letter to the network disk on the computer you are managing. This was the first available drive letter, starting from the beginning of the alphabet.

Use Disk Administrator to map a persistent drive letter. The following steps describe how to do this.

13. Choose which drive letter you want to use. If possible, select a letter near the end of the alphabet.

During system startup, drive letters are automatically allocated from the start of the alphabet. When the network disk software starts up, if the drive letter you chose has already been allocated to another disk that came online earlier in the startup sequence, the first available drive letter is allocated to the network disk. The network disk software starts late in the startup sequence, so select a letter near the end of the alphabet to reduce the probability that the drive letter is allocated to another disk.

In a cluster, use the NDMgr or SnapMgr DRIVES command (page B-2) to find out which drive letters are available on every node in the cluster. Choose one of these drive letters.

14. Use Disk Administrator to map the drive letter to the network disk and make it persistent.

In Disk Administrator, select the network disk and deassign its current drive letter, even if it already has the same drive letter as the one you chose in Step 13.

Now assign the drive letter you chose in Step 13 to the network disk. The drive letter is now persistent on the local computer.

15. In a cluster, repeat Step 14 on every node in the cluster.

Use Cluster Administrator to move the group that contains the SCE Connected Disk resource to each node, then run Disk Administrator on the node to map the drive letter and make it persistent.

In a cluster

The Network Disk Connection wizard automatically creates a cluster group and a cluster resource in that group. Both the group and the resource have the name `\\servername\netdiskname`, where:

- *servername* is the name of the standalone computer or cluster that is serving the disk
- *netdiskname* is the name of the network disk

The new resource is of type SCE Connected Disk. The wizard brings the new group and resource online.

7.4 Changing Settings

The following sections describe how to change the IP subnet configuration and how to change the client system for a served network disk.

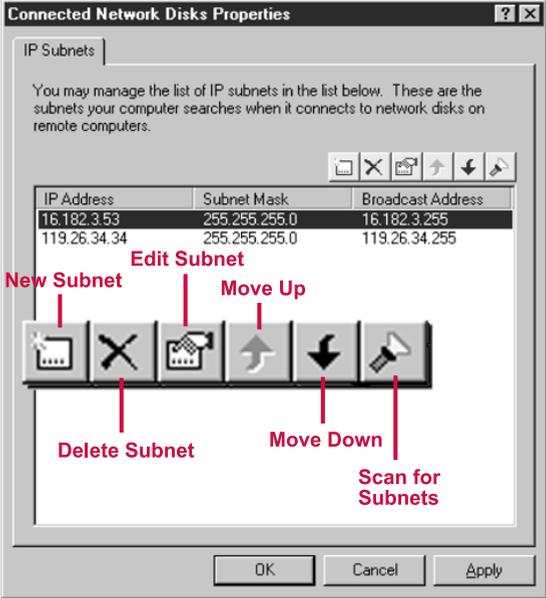
7.4.1 Changing the IP Subnet Configuration

When you install the network disk client software, the installation procedure automatically creates a default IP subnet configuration. It scans the network cards fitted to your computer. This scan creates a default configuration that lets your computer connect to network disks being served by computers and clusters in its own IP subnet.

To check the current configuration, you can either:

- Double-click the **Network Disks** icon in the **Control Panel**.
- Use the Network Disk Manager snap-in. Select the **Connected Network Disks** folder, then right-click **Properties**.

You see the IP Subnets screen.



You may want to change the IP subnet configuration in the following situations:

- If you want to connect to network disks being served by computers or clusters in another subnet, use the **New Subnet** toolbar button on the IP Subnets screen.

- If you have added a new network card and want to be able to use it to connect to network disks, use the **Scan for Subnets** toolbar button on the IP Subnets screen. This scans the network cards fitted to the computer and adds an entry for each card that is configured for TCP/IP.
- If you need to change information about a subnet, use the **Edit Subnet** toolbar button.
- To improve performance by changing the order of the subnets in the list. If your computer has two or more network cards configured for TCP/IP, make sure that it uses the fastest link when it connects to served network disks.

To do this, make sure that the IP subnet associated with the fastest card is at the top of the list. For example, if you have an Ethernet card and a Gigabit Ethernet card, make sure the IP subnet associated with the Gigabit Ethernet card is listed first.

Use the **Move Up** and **Move Down** toolbar buttons to change the order of the entries in the list.

▲ **Tip:** The changes do not take effect immediately. The next time the computer connects or reconnects to a network disk, it will use the new search order. The new order does not affect any existing network disk connections.

- If you want to use dedicated network links for network disks, use the **Delete Subnet** toolbar button to remove the entries for the network links that you do not want to use for network disk traffic.

7.4.2 Changing the Client System for a Served Network Disk

Only one standalone computer or cluster, called the **client system**, can connect to a served network disk at any point in time. If you want to change the client system, disconnect the current client system from the network disk **before** you connect the new client system.

7.5 Using a Served Network Disk as the Cluster Quorum Disk

It is possible to use a served network disk as a cluster quorum disk. This allows you to create a cluster that uses the network as its storage

interconnect and that does not have a shared storage bus. Its quorum disk is actually a disk on another cluster that is served over the network.

- ▲ **Tip:** The system serving the network disk should be a cluster, not a standalone computer. Use a reliable, preferably dedicated, network link between the cluster serving the network disk and the cluster using it as a quorum disk.

7.5.1 Steps When the Cluster Already Exists

To set up a served network disk as the cluster quorum disk, follow these steps if the cluster already exists:

1. Connect to the served network disk that you want to use as the cluster's quorum disk. Follow the steps in *Connecting to a Served Network Disk* on Page 7-8.
2. Make the drive letter mapped to the network disk persistent throughout the cluster. Steps 4 to 7 describe how to do this.
3. Use the NDMgr or SnapMgr DRIVES command (page B-2) to find out which drive letters are available on every node in the cluster. Choose which of these drive letters you want to use.
4. Run Disk Administrator on the node on which the network disk is currently online.

Select the network disk and deassign its current drive letter, even if it already has the same drive letter as the one you chose in Step 4.

Now assign the drive letter you chose in Step 4 to the network disk. The drive letter is now persistent on the current node.

5. Use Cluster Administrator to move the connected disk resource to another node in the cluster. Repeat Step 5 on the new node to make the drive letter persistent on that node.
6. Repeat Step 6 until the drive letter is persistent throughout the cluster.
7. Using Cluster Administrator, check the possible owners of the connected disk resource. Make sure that all the nodes in the cluster are listed as its possible owners.
8. Still using Cluster Administrator, change the quorum resource from a disk on a shared storage bus to the connected disk resource, as follows: Select the name of the cluster, then right-click **Properties**. Click the **Quorum** tab. In the Quorum resource list box, select the connected disk resource, then click **OK**.

The network disk is now the cluster's quorum disk.

7.5.2 Steps If the Cluster Does Not Already Exist

To create a cluster that does not have a shared storage bus:

1. On the first node on which you install the cluster software, use the undocumented `localquorum` switch on the MSCS setup command. For example:

```
C:\> D:  
D:\> CD \MSCS\CLUSTER\I386  
D:\MSCS\CLUSTER\I386> setup -localquorum
```

Answer the installation questions as appropriate, choosing to create a new cluster. This creates the quorum log file on the local system disk.

2. On the other nodes, install the cluster software in the normal way. For example:

```
C:\> D:  
D:\> CD \MSCS\CLUSTER\I386  
D:\MSCS\CLUSTER\I386> setup
```

Answer the installation questions as appropriate, choosing to join the existing cluster.

3. Follow Steps 1 to 9 of Steps When the Cluster Already Exists on Page 7-15.

Do not delete the Local Quorum resource. You may need this resource in the future, for example, to reconfigure the MSCS cluster software.

7.6 Tuning Tips

Here are some tips to optimize the performance of network disks:

- If the client computer has two or more network cards configured for TCP/IP, make sure that it uses the fastest link when it connects to served network disks.

To do this, double-click the **Network Disks** icon on the client computer's **Control Panel**. This displays a list of IP subnets. Make sure that the IP subnet associated with the fastest card is at the top of the list. If necessary, use the **Move Up** and **Move Down** toolbar buttons to change the order of the entries in the list.

- Give the client computer more memory.

The performance of network disks is sensitive to the amount of memory available to the client computer and to the workstations that access the file shares on that computer. More memory results in a bigger system cache, which means fewer I/Os are issued to disk.

Check the I/O statistics on the computer serving the network disk. Select the **Network Disk Manager** snap-in for the server. Click the **Served Network Disks** folder, then right-click **Properties**. Click the **Statistics** tab to check the ratio of reads and writes to disk at the server computer.

If it shows a high proportion of reads, it could mean that the cache of the client computer is not big enough. Increase the size of the client cache to reduce the number of I/Os that are sent over the network.

- Upgrade the network link between the client and server computers. For example, upgrade it from Ethernet to Gigabit Ethernet. The heavier the I/O load to disk, the faster network link you want.
- Use a dedicated network link. If your network is already heavily loaded, consider setting up a dedicated network link for your network disks.

Double click the **Network Disks** icon on the client computer's **Control Panel**. On the IP Subnets screen, use the **Delete Subnet** toolbar button to remove all the entries except those for the network cards that you want to use exclusively for network disk traffic.

7.7 Showing Information

The following sections describe how to show information about served network disks and connected network disks.

7.7.1 Showing Information about Served Network Disks

To find out whether a disk is being served as a network disk:

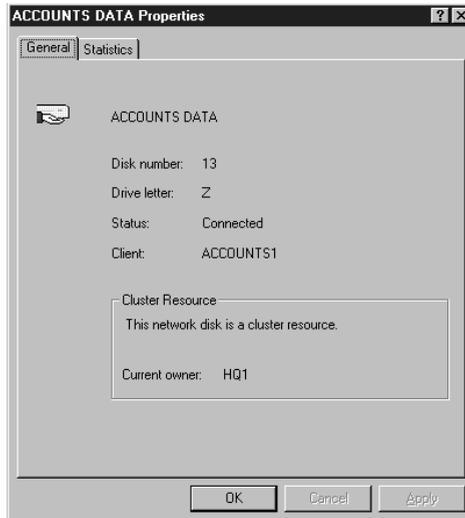
1. Run the MMC Management Console.
Click **Start**, then **Programs**, then **Compaq StorageWorks Virtual Replicator**, and then **Management Console**.
2. In the left-hand Scope pane, select the **Network Disk Manager** snap-in for the standalone computer or cluster that you want to check.
If the snap-in has not yet been added to your console, add it now. For instructions on how to do this, see "Adding a Snap-In" on Page 5-4.

3. In the left-hand Scope pane, click the **Served Network Disks** folder.

The right-hand Results pane shows all the networks disks that the standalone computer or cluster is serving. In a cluster, it shows only disks associated with served disk resources that are currently online. For each disk it shows the following information.

Column	Description
Name	The name of the network disk.
Disk	The disk number. In a cluster, this is the disk number on the node that currently owns the SCE Served Disk resource associated with the disk. The disk number may change when you restart the computer, or when the served disk resource fails over in a cluster.
Client	The name of the computer that is currently connected to the disk. If a cluster is connected to the disk, it shows the name of the node that currently owns the SCE Connected Disk resource associated with the disk.
Owner	The name of the cluster node that currently owns the SCE Served Disk resource associated with the disk. You do not see this column on a standalone computer.

For more information, select the disk then right-click **Properties**.



The **General** tab shows the drive letter that is mapped to the disk. The **Statistics** tab shows the following counters.

Counter	Description
Read count	The number of read I/Os issued to the served disk.
Data read	The amount of data read from the served disk.
Write count	The number of write I/Os issued to the served disk.
Data written	The amount of data written to the served disk.
Successes	The number of successful attempts to connect (or reconnect) to the served disk.
Failures	The number of failed attempts to connect to the served disk. An attempt to connect fails if another computer is already connected to the disk, or if the remote computer is not on the list of registered clients.

The counters are zeroed when a standalone computer restarts.

In a cluster:

- The counters are zeroed when the disk comes online on a node (for example, after cluster failover) if the served disk is a virtual disk or snapshot.
- If the served disk is **not** a virtual disk or snapshot, the counters are zeroed when the cluster restarts, and they measure cumulative information for that particular node since cluster startup.

7.7.2 Showing Information about Connected Network Disks

To find out whether a disk on a computer or cluster is a network disk that is being served by a remote computer or cluster:

1. Run the MMC Management Console.
 Click **Start**, then **Programs**, then **Compaq StorageWorks Virtual Replicator**, and then **Management Console**.
2. In the left-hand Scope pane, select the **Network Disk Manager** snap-in for the standalone computer or cluster that you want to check.
 If the snap-in has not yet been added to your console, add it now. For instructions on how to do this, see “Adding a Snap-In” on Page 5-4.

3. In the left-hand Scope pane, click the **Connected Network Disks** folder.

The right-hand Results pane shows all the networks disks that the standalone computer or cluster is connected to. In a cluster, it shows only network disks whose SCE Connected Disk resources are currently online.

For each disk it shows the following information.

Column	Description
Name	\\ <i>servername</i> \ <i>netdiskname</i> , where <i>servername</i> is the name of the standalone computer or cluster that is serving the disk, and <i>netdiskname</i> is its network disk name.
Disk	The disk number. In a cluster, this is the disk number on the node that currently owns the SCE Connected Disk resource associated with the disk. The disk number may change when you restart the computer, or when the connected disk resource fails over in a cluster.
Drive Letter	The drive letter currently mapped to the disk.
Owner	The name of the cluster node that currently owns the SCE Connected Disk resource associated with the disk. You do not see this column on a standalone computer.

An icon to the left of the name shows the status. For example, it has a warning triangle if the connection has been lost.

Show the status by selecting the disk then right-clicking **Properties**.

Status	Description
Connecting	In the process of connecting.
Connected	Connected.
Disconnecting	In the process of disconnecting.
Disconnected	Disconnected.
Failed	Failed to connect at last startup, because either the network was down or the server was unavailable at the time; will try to connect again at next startup. You only see this status on a standalone computer.
Reconnecting	Currently trying to reconnect; the client lost the connection. The server or the network link failed, or the disk failed over within the server cluster. In a cluster, the node that currently owns the connected disk resource keeps trying to reconnect until the link can be re-established. If it lost the connection because the network link failed, the failure may not affect every node in the cluster. If another node can still see the system serving the network disk, use Cluster Administrator to move the connected disk resource to that node.

7.8 Disconnecting from a Served Network Disk

Disconnecting from a network disk generates a data cache flush on a local system. All the data that has not yet been written to the disk is flushed to disk before it is disconnected.

If the network connection has been lost, the client status of the network disk displays the message, *reconnecting*:

- On a standalone computer, the attempt to disconnect will hang, if necessary forever, until the computer reconnects and writes its cached data to disk.
- In a cluster, the software continues trying to reconnect for 30 seconds. If the connection has not been re-established within 30 seconds, the disk is disconnected and data in the cache that had not been written to disk is lost.

Steps

To disconnect a standalone computer or cluster from a served network disk:

1. In a cluster, make sure that no resources depend on the network disk you want to disconnect. You cannot disconnect from the disk if any cluster resources are dependent on the SCE Connected Disk resource. Use Cluster Administrator to remove the dependencies on the SCE Connected Disk resource.
2. Run the MMC Management Console.
Click **Start**, then **Programs**, then **Compaq StorageWorks Virtual Replicator**, and then **Management Console**.
3. In the left-hand Scope pane, select the **Network Disk Manager** snap-in for the computer or cluster that you want to disconnect from the network disk.
If the snap-in has not yet been added to your console, add it now. For instructions on how to do this, see “Adding a Snap-In” on Page 5-4.
4. In the left-hand Scope pane, select the **Connected Network Disks** folder.
5. In the right-hand Results pane, select the network disk you want to disconnect, then right-click **Delete**.
6. Click **Yes** to confirm that you want to delete the connection.

In a cluster

Compaq StorageWorks Virtual Replicator automatically deletes the cluster resource of type SCE Connected Disk that was created when you connected to the disk. It also deletes the group that the resource was in, if that group is now empty.

7.9 Stopping the Serving of a Network Disk

When you stop serving a network disk, the disk comes back online on your computer. You can then access the disk on the local machine, for example, by double-clicking its icon in Windows NT Explorer.

- ▲ **Tip:** If the disk is a virtual disk or snapshot, use the Snapshot Manager snap-in to stop serving it. Select the disk, and then right-click **Task** then **Stop serving**.
- ▲ **Tip:** You cannot stop serving a network disk if any node in the cluster is down.

Steps

To stop serving a network disk:

1. Run the MMC Management Console.
Click **Start**, then **Programs**, then **Compaq StorageWorks Virtual Replicator**, and then **Management Console**.
2. In the left-hand Scope pane, select the **Network Disk Manager** snap-in for the standalone computer or cluster that is serving the network disk.
If the snap-in has not yet been added to your console, add it now. For instructions on how to do this, see “Adding a Snap-In” on Page 5-4.
3. In the left-hand Scope pane, select the **Served Network Disks** folder.
4. In the right-hand Results pane, check that the Client column is empty for the network disk you want to stop serving.
If the Client column is not empty, disconnect the client from the network disk. For instructions on how to do this, see “Disconnecting from a Served Network Disk” on Page 7-21.

- ▲ **Tip:** If a cluster is connected to the network disk, the Client column is empty whenever the connected disk resource is failing over from one node in that cluster to another.

In the right-hand Results pane, select the network disk you want to stop serving, and then right-click **Task** then **Stop serving**.

5. Click **Yes** to confirm that you want to stop serving the disk.

In a cluster

Compaq StorageWorks Virtual Replicator automatically deletes the cluster resource of type SCE Served Disk that was created when you served the disk. It also deletes the group that the resource was in, if that group is now empty.

8 *Compaq Batch Scheduler Wizards*

Compaq StorageWorks™ Virtual Replicator provides six *Compaq Batch Scheduler* (CBS) Task Automation wizards that allow you to schedule certain tasks that will be carried out using the *Compaq Batch Scheduler*. The Task Automation wizards provide a quick way to create a schedule for various snapshot activities.

Once you have created a virtual disk, you have access to these wizards. Using these wizards requires you to log on to the *Compaq Batch Scheduler APS*; therefore, you must have a valid CBS user name and password. Also, while the wizards are active, the CBS Coordinator or Monitor must not be running at the same time on that system.

Note: *Compaq Batch Scheduler Sentinel* must be installed on the node on which you are running *Compaq StorageWorks Virtual Replicator*.

This chapter describes the following topics:

- Task Automation wizards (Page 8-1)
- Using the Batch Scheduler wizards (Page 8-2)
- Modifying the Task Automation Schedules (Page 8-3)

For information on how to use these wizards, refer to the *Compaq Batch Scheduler* wizards online help.

8.1 Task Automation Wizards

StorageWorks Virtual Replicator provides the following Task Automation wizards:

- Create Snapshot

The Create Snapshot wizard schedules the creation of a snapshot of a selected virtual disk at a specified time and frequency. The operation includes deleting any previous snapshot by the same name. The snapshot creation is scheduled using the *Compaq Batch Scheduler*.

■ Delete Snapshot

The Delete Snapshot wizard schedules the deletion of a snapshot with the name you specify at a predetermined time. If the snapshot name provided does not exist at the time the scheduler runs, no action is taken. The snapshot deletion is scheduled using the *Compaq Batch Scheduler*.

■ Restore From Snapshot

The Restore From Snapshot wizard schedules the re-creation of a virtual disk from a snapshot. If specified, upon successful completion of the schedule, the previous virtual disk and snapshot are both deleted. The Restore operation is scheduled using the *Compaq Batch Scheduler*.

■ Snapshot For Backup

The Snapshot For Backup wizard schedules a backup operation of a snapshot to tape. The snapshot and backup operations are scheduled using the *Compaq Batch Scheduler*.

■ Snapshot Watchdog

The Snapshot Watchdog wizard schedules a job, at a given frequency, which monitors the resource utilization of a snapshot. The snapshot can be optionally deleted if the resource utilization exceeds the present values. The snapshot monitoring and deletion are scheduled using the *Compaq Batch Scheduler*.

■ Workday Snapshot

The Workday Snapshot wizard schedules the creation of a snapshot at a defined time (for example, start of the work day) and delete the snapshot after the defined time (for example, end of the work day). The snapshot creation and deletion are scheduled using the *Compaq Batch Scheduler*.

8.1.1 Using the Batch Scheduler Wizards

To create a schedule with the *Compaq Batch Scheduler* wizards, right click on the appropriate virtual disk or snapshot. Select the **All Tasks** menu item, and then select **Compaq Batch Wizards**. You will see the list of available wizards.

Once you select and run a wizard, you are asked to login to the *Compaq Batch Scheduler* APS. You will need a valid **Compaq Batch Scheduler** user name and password to continue with the wizard.

Once you have completed the information required by the wizard, you are then notified that your schedule is about to be entered into the *Compaq*

Batch Scheduler. At this point, you have the option of launching the schedule immediately or saving it to launch at another time.

To modify the schedule, use the *Compaq Batch Scheduler Coordinator* or the *Compaq Batch Scheduler Monitor*. See the next section for this information.

8.1.2 Modifying Task Automation Schedules

Once you have created a schedule for a node using a Task Automation wizard, use the *Compaq Batch Scheduler Coordinator* or *Monitor* to modify it. You can also duplicate the schedule to run on another node or a group of nodes.

To change the schedule so that it runs on more than one node, make the following changes to every process:

- Change the name of the machine(s) on which you want the schedule to run, adding the new node names.
- If the command field of the process contains the Snap Manager, delete the current node name in the parameter field.

These changes must be made for each process in the schedule.

- ▲ **Note:** The original schedule was created for a single node, with known snapshots, virtual disks, and drive letters. If you are adding nodes to the schedule, take into account the availability of snapshots, virtual disks, and drive letters for those nodes.

See the documentation for the *Compaq Batch Scheduler* for information on using the *Coordinator* and *Monitor* to modify the schedules.

9 Troubleshooting

This chapter lists possible problems that may occur and answers to help you solve them. The following problems are described in this chapter:

- Applications using a network disk fails or hangs (Page 9-1)
- Lost Delayed Write Errors (Page 9-4)
- Applications Return Failed Write Errors (Page 9-5)
- Accidentally Delete a Virtual Disk (Page 9-6)
- Snap-in and Commands Do Not Show All Pools in a Cluster (Page 9-7)
- Serving Tasks Missing from the Snapshot Manager Snap-in (Page 9-7)
- Disk Structure Corrupt Errors when Accessing a Snapshot (Page 9-8)
- Disk Structure Corrupt Errors when you Stop Serving a Disk (Page 9-9)
- Pool Free Space Fell by More Than the Amount of Data Written (Page 9-10)
- Cannot Connect to a Network Disk in Another Subnet (Page 9-10)
- SCE Connected Disk Resource Fails to Come Online in a Cluster (Page 9-11)
- Reformatting a Virtual Disk Does Not Work (Page 9-11)
- Performance Monitor shows zeros for network disks (Page 9-12)
- Reconstructing Pool (Page 9-12)
- Unable to Delete a Snapshot or a Virtual Disk (page 9-12)

9.1 Application using a Network Disk Fails or Hangs

If an application using a network disk fails or hangs, check whether you have lost the connection to the network disk. Follow these steps:

1. Run the MMC Management Console.

Click **Start**, then **Programs**, then **Compaq StorageWorks Virtual Replicator**, and then **Management Console**.

2. In the left-hand Scope pane select the **Network Disk Manager** snap-in for the standalone computer or cluster on which the application is running.

In the left-hand Scope pane, select the **Connected Network Disks** folder.

3. If you can see the connected network disk in the right-hand Results pane, go to Step 4.

If you are unable to see the connected network disk in the right-hand Results pane, the action you take depends on whether the client system is a standalone computer or cluster:

On a standalone computer

Re-connect it to the network disk again, following the instructions in the section “Connecting to a Served Network Disk” on Page 7-8.

In a cluster

Use Cluster Administrator to check whether there is an SCE Connected Disk resource for the network disk.

If this resource does not exist, re-connect the cluster to the network disk, following the instructions in “Connecting to a Served Network Disk” on Page 7-8.

If the SCE Connected Disk resource does exist, try to bring it online, then go to Step 4.

If you cannot bring the resource online, check in the application event log to find out why it is not online. Look for entries that have SCE Connected Disk in the Source column.

Double-click the entry to show the Description text. It will tell you why the resource failed to come online, for example, it may say that the problem was caused by a failure in the network link.

4. Select the connected network disk, then right-click **Properties**.

The Status line tells you whether you have a connection to the network disk.

Make a note of the status, then click **OK** to dismiss the Properties screen.

If the status is **Connected**, the problem is not caused by a lost connection to the network disk.

If the status is anything other than **Connected**, go to Step 5.

5. In the left-hand Scope pane select the **Network Disk Manager** snap-in for the computer or cluster that is serving the network disk.

In the left-hand Scope pane, select the **Served Network Disks** folder.

6. Can you see the served network disk in the right-hand Results pane?

Yes	<p>Select the served network disk, then right-click Properties. Check the Status shown on the General tab.</p> <p>If the status is not Configured, go to Step 7.</p> <p>If the status is Configured, the system failed to serve the network disk at last startup because the disk was not available at that time. Follow these steps:</p> <p>Disconnect the client system from the network disk. If the client is a standalone computer, you will need to restart the computer before you can disconnect it from the network disk. Then check whether someone has removed the disk; if necessary, replace the disk. Make sure that the disk's cable connections are tight. Now restart the computer, and check that the disk is being served. Finally connect the client to the network disk. Depending on what happened to the disk, you may have to do some level of repair before you can use it – for example, you may need to run <code>chkdsk</code> on it.</p>
No	<p>Use Cluster Administrator to check whether there is an SCE Served Disk resource for the disk.</p> <p>If this resource does not exist, someone could have accidentally stopped serving the disk. Serve the disk again, following the instructions in "Serving a Disk" on Page 7-1.</p> <p>If the SCE Served Disk resource does exist, try to bring it online, then go to Step 8. If you cannot bring it online, go to Step 7.</p>

7. Check that the network link is OK between the network disk's client and server systems.

Use the ping command to check the network link.

If either the client or server system is a cluster, use the ping command to check all the connections between all the nodes. If one node has a network connection and the other does not, use Cluster Administrator to move either the SCE Connected Disk or SCE Served Disk resource (depending on whether the cluster is the client or the server system) to the node that has the network connection, then go to Step 8.

If necessary, repair the network link, then go to Step 8.

8. You should have now fixed the problem that caused the connection to the network disk to be lost. The connection is automatically re-established unless the status of the connected network disk in Step 4 was Failed.

If the status in Step 4 was Failed, go to Step 9 to re-establish the connection to the network disk.

9. To re-establish the connection when the status was Failed, you need to delete the connection and then create a new one. This step describes how to do this.

Use the **Network Disk Manager** snap-in you used in Step 2.

In the left-hand Scope pane, select the **Connected Network Disks** folder.

In the right-hand Results pane, select the disk, right-click **Delete**, then click **Yes** to confirm that you want to delete the connection.

In the left-hand Scope pane, select the **Connected Network Disks** folder, then right-click **New** then **Connected Network Disk**. Follow the instructions in the wizard.

9.2 Lost Delayed Write Errors

- ▲ **Tip:** In a cluster, you may also get lost delayed write errors when you move an SCE Connected Disk resource from one node to another **and** the cluster has lost the connection to the network disk.

If you are using snapshots and you get popup windows that say that delayed writes have been lost, follow these steps to check whether you have a full pool:

1. Run the MMC Management Console.
Click **Start**, then **Programs**, then **Compaq StorageWorks Virtual Replicator**, and then **Management Console**.
2. In the left-hand Scope pane select the **Snapshot Manager** snap-in for the standalone computer or cluster that is getting the popup errors.

3. With the snap-in highlighted, check the **Free** column in the right-hand Results pane. It shows the free space of all the pools on that computer or cluster.

Do any pools have 0 MB in the **Free** column?

Yes	Go to Step 4.
No	The symptoms you are seeing are not caused by a full pool. Pursue other lines of investigation to diagnose the problem.

4. If a pool has 0 MB in the **Free** column, select the pool in the left-hand Scope pane, and expand it to display its virtual disks and snapshots.

Are there any snapshots in the pool?

Yes	Go to Step 5.
No	The symptoms you are seeing are not caused by a full pool. Pursue other lines of investigation to diagnose the problem.

5. Take immediate action to either add a storage unit to the pool, or delete a snapshot or virtual disk from the pool.

▲ **Tip:** Deleting files will not increase the pool free space.

To add a storage unit to the pool, select the pool, then right-click **Properties**. Click the **Storage Units** tab, then click **Add**. Select a storage unit then click **Add**, **OK**, then **OK**. For more information, see “Adding Storage Units to a Pool” on Page 6-12.

To delete a snapshot or virtual disk, select it then right-click **Delete**. Click **OK** to confirm that you want to delete the disk and erase all the data on it.

9.3 Applications Return Failed Write Errors

If you are using snapshots and applications return errors when they try to write data to disk, follow these steps to check whether you have a full pool:

1. Run the MMC Management Console.

Click **Start**, then **Programs**, then **Compaq StorageWorks Virtual Replicator**, and then **Management Console**.

2. In the left-hand Scope pane select the **Snapshot Manager** snap-in for the standalone computer or cluster that is getting the failed write errors.

3. With the snap-in highlighted, check the **Free** column in the right-hand Results pane. It shows the free space of all the pools on that computer or cluster.

Do any pools have 0 MB in the **Free** column?

Yes	Go to Step 4.
No	The symptoms you are seeing are not caused by a full pool. Pursue other lines of investigation to diagnose the problem

4. If a pool has 0 MB in the **Free** column, select the pool in the left-hand Scope pane, and expand it to display its virtual disks and snapshots.

Are there any snapshots in the pool?

Yes	Go to Step 5.
No	The symptoms you are seeing are not caused by a full pool. Pursue other lines of investigation to diagnose the problem.

5. Take immediate action to either add a storage unit to the pool, or delete a snapshot or virtual disk from the pool.

▲ Deleting files does not increase the pool free space.

To add a storage unit to the pool, select the pool, then right-click **Properties**. Click the **Storage Units** tab, then click **Add**. Select a storage unit then click **Add**, **OK**, then **OK**. For more information, see “Adding Storage Units to a Pool” on Page 6-12.

To delete a snapshot or virtual disk, select it then right-click **Delete**. Click **OK** to confirm that you want to delete the disk and erase all the data on it.

9.4 Accidentally Delete a Virtual Disk

Deleting a virtual disk does not overwrite the physical disk blocks used to store its data. If you accidentally deleted a virtual disk, follow these steps **immediately** to try to salvage the data on the disk:

1. Create a new virtual disk of **exactly** the same capacity in the **same** pool.

In the New Virtual Disk wizard, map a drive letter to the new virtual disk but do not format it. Click **No** when the wizard asks if you would like to format the disk now.

2. Double-click the virtual disk's icon in Windows NT Explorer. You may see all the folders and files of the virtual disk you accidentally deleted. The new virtual disk is using exactly the same disk blocks as the one you accidentally deleted.
3. In Step 2 when you double-click the disk icon, if you see a popup that says the drive is inaccessible and does not contain a file system, you must restore the data from your backup tapes.

9.5 Snap-in and Commands Do Not Show All the Pools in a Cluster

In a cluster, the Snapshot Manager snap-in and commands show only information about pools that are online. If a pool is missing from the display, use Cluster Administrator to check if it is offline. Try to bring it online, using Cluster Administrator.

If that fails, look in the system event log to find out why the pool is not online.

Use Windows NT Event Viewer. Look for entries that have **sddriver** in the **Source** column.

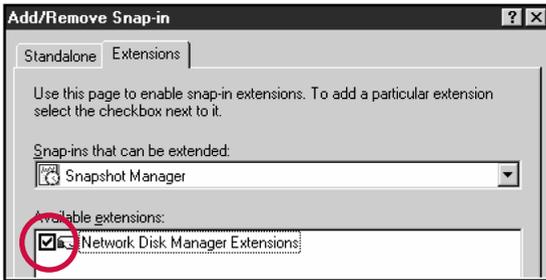
Double-click the entry to show the Description text. It may tell you why the pool failed to come online.

9.6 Serving Tasks Missing from the Snapshot Manager Snap-in

If you select a virtual disk or snapshot and there is no **Serve as Network Disk** option on MMC's **Task** menu, you need to add the Network Disk Manager Extension snap-in to the Snapshot Manager snap-in. Follow these steps:

1. On the MMC **Console** menu, click **Add/Remove Snap-in**. You see the Add/Remove Snap-in screen.
2. In the top pull-down list box, select the folder that contains the Snapshot Manager snap-in. You see a list of all the snap-ins in that folder.
3. Select the Snapshot Manager snap-in, then click the **Extensions** tab.

4. Check the square box to the left of **Network Disk Manager Extensions**, as shown in this picture.



5. Click **OK**.

9.7 Disk Structure Corrupt Errors when Accessing a Snapshot

When you access a snapshot, there are two situations where you may get a popup that says that the snapshot disk is inaccessible and that its disk structure is corrupt and non-readable:

- If the parent disk is being served as a network disk **and** you did not disconnect the remote client before you created the snapshot.

When you access the snapshot you may get disk structure corrupt errors because creating the snapshot does not flush the system cache on the remote client. As a result, the snapshot disk may be internally inconsistent because it does not include updates held in the remote client's cache. To make the snapshot disk consistent, run `chkdsk` on it.

This problem is more likely to occur with high workloads.

- ▲ **Tip:** When you serve a virtual disk as a network disk, if you want to make sure that your snapshot is internally consistent, disconnect the remote client before you create the snapshot, as described in "Using Snapshots to do Backups" on Page 6-16.

- If the architecture of the computer accessing the snapshot is different from that of the architecture of the computer that was writing to the parent disk when the snapshot was created.

When you access the snapshot you may get disk structure corrupt errors because the NTFS file system's recovery log has a slightly different format on Intel and Alpha computers.

For example, you may receive disk structure corrupt errors when you access a snapshot if:

- You have a pool on an Alpha computer and an application on that computer is accessing a virtual disk in the pool. You then create a snapshot of the virtual disk and serve it to an Intel computer.
- You serve a virtual disk to an Alpha cluster and you serve its snapshot to an Intel cluster.
- You have a pool on an Alpha computer, and you serve one of the virtual disks in the pool to an Intel computer. You then create a snapshot of the virtual disk and an application running on the Alpha computer tries to access the snapshot.

In these cases, the disk structure is not corrupt. To fix the problem, make sure that the architectures of the computers accessing the snapshot and its parent disk are the same. They must both be Intel or both Alpha.

9.8 Disk Structure Corrupt Errors when you Stop Serving a Disk

When you stop serving a network disk, the disk should come back online on your computer. You should be able to access it directly, for example, by double-clicking its icon in Windows NT Explorer.

When you try to access the disk, if you get a popup that says that it is inaccessible and the disk structure is corrupt and non-readable, check the following:

- Does your computer have a different architecture from the client computer that was previously connected to the disk – is one Intel and the other Alpha?
- Did the client computer fail before you stopped serving the disk?
- Did the client computer not manage to reconnect before you stopped serving the disk?

If the answer to all of these questions is **yes**, attempts to access the disk on your computer are failing with file system corrupt errors because the NTFS file system's recovery log has a different format on Intel and Alpha.

NTFS needs to perform recovery because the client failed. The recovery log has the wrong format, so your attempts to access the disk fail with file system corrupt errors.

To fix this problem, follow these steps:

1. Serve the disk on your computer again.
2. Restart the client computer. It automatically reconnects to the disk and performs file system recovery.

If you cannot restart the client computer, for example, because it was destroyed in a fire, choose any computer that has the same architecture as the original client, and configure it to connect to the disk.

3. On the client computer, try to access the disk. Double-click its icon in Windows NT Explorer. If you see all the folders and files, the problem is now fixed.
4. Disconnect the client computer from the disk.
5. Stop serving the disk on your computer. Your computer should now be able to access the disk.

9.9 Pool Free Space Fell by More Than the Amount of Data Written

When a disk has a snapshot and you write data to either the snapshot or its parent disk, the pool space may fall by more than the amount of data that you write. This is because the data is written in 32 KB chunks, called segments.

When a segment contains any data modified by the write, the whole segment is copied out. For example, you write 2 KB of data and one KB is in one segment and the other is in another segment; the pool free space falls by 64 KB. Normally, the disk space used by a write is very similar to the size of the write. You should see a difference only in unusual situations, for example, when running a test program that writes small amounts of data randomly over the surface of the disk.

9.10 Cannot Connect to a Network Disk in Another Subnet

If you cannot connect to a network disk that is being served by a computer or cluster in another subnet, it could be that your routers are not configured to allow messages to be broadcast to other subnets. In this case:

1. On the **Control Panel** of your computer, double-click **Network Disks**.
2. You see the IP Subnets screen. Check that there is an entry for the computer serving the network disk. The entry should give the IP address of the computer and have a subnet mask of 255.255.255.255
If the network disk is being served by a cluster, there should be an entry for each node in the cluster.
3. If necessary, use the **New Subnet** toolbar button to add the appropriate entries to the list.

9.11 SCE Connected Disk Resource Fails to Come Online in a Cluster

If a cluster is connected to a network disk, and the SCE Connected Disk resource fails to come online, try using Cluster Administrator to manually bring it online.

If this fails, check if there is any information in the event log that can help you diagnose the problem. Use Windows NT Event Viewer. Look for entries that have **SCE Connected Disk** in the **Source** column.

Double-click the entry to show the Description text. It may tell you why the resource failed to come online, for example, it may say that the problem was caused by a failure in the network link.

9.12 Reformatting a Virtual Disk Does Not Work

When you use the Snapshot Manager snap-in or command to reformat a virtual disk, you may see the **Format Complete** popup or the **Succeeded** message text but the disk is not formatted.

This happens if any application is accessing the disk. For example, it happens if you are running Windows NT Explorer, or if you have a Windows NT command window and the command prompt includes the drive letter of the virtual disk you are reformatting.

To fix this problem, stop all applications from accessing the virtual disk before you use the Snapshot manager snap-in or command to reformat the disk. For example, close Windows NT Explorer, or change the prompt in the command window to another drive, then use the snap-in to reformat the disk again

9.13 Performance Monitor Shows Zero Network Disks

In this release of *Compaq StorageWorks Virtual Replicator*, you cannot use the Windows NT Performance Monitor to monitor network disks.

Performance Monitor displays zeros for every counter when you use it to monitor network disks.

9.14 Reconstructing a Pool

If for any reason, you need to reconstruct a pool, follow these steps:

1. Decide which storage units you want to use in the new pool.
2. Create a new pool using the storage units. Follow the instructions in "Creating Pools, Virtual Disks, and Snapshots" on Page 6-2.
3. Create virtual disks in the new pool. Follow the instructions in "Creating a Virtual Disk" on Page 6-5.
4. Use your backup tapes to restore saved data to the new virtual disks.

9.15 Unable to Delete a Snapshot or a Virtual Disk

You cannot delete a snapshot or a virtual disk if it is in use. Snapshots and virtual disks can be in use if an application has files open or services have connections to the snapshots or virtual disks..

If you cannot delete a snapshot or virtual disk, do the following:

- Close any open files on the virtual disk or snapshot.
- Shut down any applications with open files on the virtual disk or snapshot.
- Stop services that have connections to the virtual disk or snapshot, for example, diskperf (a component of Perfmon) or virus checking software.

A SnapMgr Commands

This appendix describes the commands that you can use to manage pools, and the virtual disks and snapshots in them.

The following table gives a summary of the commands. For information on how to use them, see “Using the Command Line Interface” on Page 5-6.

Command	Description
DRIVES	Shows which drive letters are available for you to map to virtual disks, snapshots and network disks.
MANAGE	Controls which computer or cluster you manage (only available at the SnapMgr or NDMgr prompt).
POOL	Manages pools.
SNAPSHOT	Manages snapshots.
UNITS	Shows which storage units you can use to create a new pool or add to an existing pool.
VIRTUALDISK	Manages virtual disks.

A.2 DRIVES

Syntax

DRIVES [/NODE:*nodename*]

- *nodename* is the name of the cluster node whose available drive letters you want to show.

Description

Use the DRIVES command to show which drive letters are available for you to map to virtual disks, snapshots and network disks.

In a cluster, the command shows drive letters that are available only on every node in the cluster.

If one of the nodes in the cluster is down, use the /NODE switch to find out which drive letters are available on a particular node. The DRIVES command fails if one of the nodes is down and you don't use the /NODE switch.

Examples

1. This example shows which drive letters are available on every node in the HQ cluster.

```
SnapMgr> MANAGE HQ  
SnapMgr> DRIVES
```

```
Available drive letters: DGHMNPQ
```

2. This example shows which drive letters are available on the standalone computer Backup.

```
C:\> SNAPMGR Backup DRIVES
```

```
Available drive letters: DEFGHMNPQRSTVW
```

3. This example shows which drive letters are available on the local cluster node Accounts1.

```
SnapMgr> MANAGE /LOCAL  
SnapMgr> DRIVES /NODE:Accounts1
```

```
Available drive letters: GHKMNPQRST
```

A.2 MANAGE

Summary

Use MANAGE commands to control which computer or cluster you manage. You can:

- Manage a remote computer or cluster:
MANAGE *computername*
- Manage the local computer or cluster:
MANAGE /LOCAL
- Show which computer or cluster you are managing:
MANAGE

You can use MANAGE commands only at the SnapMgr or NDMgr prompt. You cannot use them at the Windows NT command prompt.

If you are entering commands at the Windows NT command prompt, each command automatically manages the local computer or cluster unless you specify the name of a remote computer or cluster after the SNAPMGR or NDMGR command prefix.

For example, these commands show which drive letters are available, first on the Accounts cluster, and then on the local cluster:

```
C:\> SNAPMGR Accounts DRIVES
C:\> SNAPMGR DRIVES
```

A.2.1 Manage a Remote Computer

Syntax

MANAGE *computername*

- *computername* is the name of the standalone computer or cluster that you want to manage.

If you want to manage a cluster, you normally specify the cluster name (which is equivalent to specifying the name of the node that currently owns the cluster name resource). But when you are creating a pool or adding a storage unit to a pool, you need to specify the name of a node in the cluster that currently owns the resource. This is because you supply the disk numbers of the storage units you want to use, and disk numbers can be different on different nodes.

Normally it does not matter which node you manage, so you use the cluster name. But when you use the UNITS command (Page A-31) to check disk numbers and the POOL command (Page A-8 and Page A-12) to create a pool or add a storage unit to a pool, make sure you specify a node name.

Description

You can use this command only at the SnapMgr or NDMgr prompt.

By default, you manage the local standalone computer or cluster. Use this command to manage a remote computer or cluster. Subsequent commands that you enter at the SnapMgr or NDMgr prompt will manage the specified standalone computer or cluster.

Examples

1. This example manages the standalone computer called Usenet. Subsequent commands that you enter at the SnapMgr prompt will manage Usenet.

```
SnapMgr> MANAGE Usenet
```

A.2.2 Manage the Local Computer

Syntax

MANAGE /LOCAL

Description

You can use this command only at the SnapMgr or NDMgr prompt.

If you have previously issued a MANAGE command to manage a remote computer or cluster, use this command to revert to managing the local computer or cluster. Subsequent commands that you enter at the SnapMgr or NDMgr prompt will manage the local standalone computer or cluster.

Examples

1. This example shows which drive letters are available, first on the Accounts cluster, and then on the local cluster.

```
SnapMgr> MANAGE Accounts  
SnapMgr> DRIVES  
SnapMgr> MANAGE /LOCAL  
SnapMgr> DRIVES
```

A.2.3 Show the Managed Computer

Syntax

MANAGE

Description

You can use this command only at the SnapMgr or NDMgr prompt.

This command shows which standalone computer or cluster you are currently managing. This is the computer or cluster that the subsequent commands that you enter at the SnapMgr or NDMgr prompt will manage.

Examples

1. In this example, you are managing the standalone computer called Usenet. Subsequent commands issued at the SnapMgr prompt will manage Usenet.

```
SnapMgr> MANAGE
Managing node USENET
```

2. In this example, you are managing the Accounts cluster. Subsequent commands issued at the SnapMgr prompt will manage the Accounts cluster.

```
SnapMgr> MANAGE
Managing cluster node ACCOUNTS1
Managing cluster ACCOUNTS
Cluster members: ACCOUNTS1      ACCOUNTS2
```

A.3 POOL

Summary

Use POOL commands to manage pools. You can:

- Create a pool:
POOL *pool* /UNITS:*disknumbers*
- Add a storage unit to a pool:
POOL *pool* /ADD:*disknumber*
- Delete a pool:
POOL *pool* /DELETE
- Show pools:
POOL
POOL *pool*
POOL *pool* /STATISTICS
POOL *pool* /UNITS

A.3.1 Create a Pool

Syntax

`POOL pool /UNITS:disknumbers`

- *pool* is the name that you want to give to the new pool. It can be up to 23 characters. If you want to use spaces in the name, enclose it in quotation marks.

You must choose a name that is different from that of any other pool, virtual disk, or snapshot on the standalone computer or cluster you are managing. In a cluster, the name must also be different from that of any existing cluster group or cluster resource.

Note that you cannot change the name of the pool later.

- *disknumbers* is a list of the disk numbers of the storage units you want to use in the pool. Use commas to separate the entries in the list. Do not use spaces between the entries in the list. Use the UNITS command (page A-31) to find out which storage units are available and what their disk numbers are.

Description

This command creates a pool from the specified storage units. It marks the storage units as offline, so that users cannot access them directly. In a cluster, the command fails if any of the nodes in the cluster is down.

You can have up to 8 storage units in a pool:

- They can be standard single-spindle disks or controller-based fault-tolerant disk sets. They can have different capacities and be from different manufacturers.
- They should all have the same read-write, redundancy and failure characteristics. For example, they should all be standard disks, or they should all be StorageWorks RAID 5 arrays, or they should all be mirror sets.

- In a cluster, use disks only on a shared storage bus; do not use local disks. Only use disks that have **Shared** in the **Type** column shown by the UNITS command (page A-31).
- Do not use network disks that are being served by remote computers or clusters.
 - ▲ **Note:** Do not use disks that are part of volume sets, mirror sets, or stripe sets created using Disk Administrator.

You can have any number of pools on a standalone computer or cluster. The number is limited only by the number of storage units that are available.

Each storage unit must be a non-removable disk that is online and does not contain any partitions. In a cluster, it must not be configured as a cluster resource of type Physical Disk. If necessary, before you create the pool, use Cluster Administrator to delete the Physical Disk resource and Disk Administrator to delete partitions.

In a cluster

If you are managing a cluster, the command automatically creates a cluster group that has the same name as the pool. Then it creates resources in the pool group:

- It creates a resource of type SCE Storage Unit for every storage unit in the pool. The name of the new resource is Storage_{N}, where N is a 32-digit hexadecimal number that uniquely identifies the storage unit.
- It creates a resource of type SCE Pool for the pool itself. This resource has the same name as the pool. It is dependent on all of the storage unit resources. A pool must always be able to access all of its storage units, so it can't come online until all of the storage units are online.

The command brings the pool group and all of these resources online.

- ▲ **Note:** Do not use Cluster Administrator to rename the pool or storage unit resources, or remove the dependency of the pool resource on the storage unit resources.

Examples

1. This example creates a pool called StorageWorks RAID Pool on the HQ cluster. It uses the UNITS command to find out what storage units are available and what their disk numbers are. Because disk numbers can be different on different nodes, it specifies the name of a node, and not the cluster name, on the MANAGE command. Then it uses the POOL command to create the pool.

```
SnapMgr> MANAGE HQ1  
SnapMgr> UNITS
```

Capacity (MB)	Disk Number	Type
4091	1	Shared
8768	3	Shared
8768	4	Shared
4091	6	Shared
2007	7	Shared

```
SnapMgr> POOL "RAID Pool" /UNITS:3,4
```

2. This example creates a pool called News Pool on the HQ cluster. It uses the UNITS command to find out which storage units are available and what their disk numbers are. Then it uses the POOL command to create the pool. Because disk numbers can be different on different nodes, it specifies the name of a node, and not the cluster name, on the UNITS and POOL commands.

```
C:\> SNAPMGR HQ1 UNITS
```

Capacity (MB)	Disk Number	Type
2007	7	Shared

```
C:\> SNAPMGR HQ1 POOL "News Pool" /UNITS:7
```

A.3.2 Add a Storage Unit to a Pool

Syntax

`POOL pool /ADD:disknumber`

- *pool* is the name of the pool to which you want to add the storage unit.
- *disknumber* is the disk number of the storage unit that you want to add. Use the UNITS command (Page A-31) to find out which storage units are available and what their disk numbers are.

Description

This command adds a storage unit to an existing pool. It marks the storage unit as offline so that users cannot access the storage unit directly. In a cluster, the command fails if any of the nodes in the cluster are down or if the pool cluster resource is offline.

You can add a storage unit to a pool when users are accessing its virtual disks and snapshots.

You can have up to 8 storage units in a pool:

- They can be standard single-spindle disks or controller-based fault-tolerant disk sets. They can have different capacities and be from different manufacturers.
- They should all have the same read-write, redundancy and failure characteristics as the existing storage units in the pool. For example, if the existing storage units are StorageWorks RAID 5 arrays, you should add only a StorageWorks RAID 5 array.
- In a cluster, use disks only on a shared storage bus; do not use local disks. Use disks that have only **Shared** in the **Type** column shown by the UNITS command (Page A-31).

- Do not use network disks that are being served by remote computers or clusters.
 - ▲ **Note:** Do not use disks that are part of volume sets, mirror sets, or stripe sets created using Disk Administrator.

The storage unit you are adding must be a non-removable disk that is online and does not contain any partitions. In a cluster, it must not be configured as a cluster resource of type Physical Disk. If necessary, before you add it to the pool, use Cluster Administrator to delete the Physical Disk resource and Disk Administrator to delete partitions.

In a cluster

If you are managing a cluster, the command automatically:

- Creates a resource of type SCE Storage Unit in the cluster group that contains the pool resource. The name of the new resource is Storage_{N}, where N is a 32-digit hexadecimal number that uniquely identifies the storage unit.
- Changes the properties of the pool resource to make it dependent on all of the new storage units resources.

The command brings the storage unit resource online.

- ▲ **Note:** You must not use Cluster Administrator to rename the storage unit resource, or remove the dependency of the pool resource on it.

Examples

1. This example adds a storage unit to the pool StorageWorks RAID Pool on the HQ cluster. It uses the UNITS command to find out what storage units are available and what their disk numbers are. Because disk numbers can be different on different nodes, it specifies the name of a node, and not the cluster name, on the MANAGE command. It uses the POOL command to add the storage unit.

```
SnapMgr> MANAGE HQ1
SnapMgr> UNITS
```

Capacity (MB)	Disk Number	Type
4091	1	Shared
4091	6	Shared

```
SnapMgr> POOL "RAID Pool" /ADD:6
```

2. This example adds a storage unit to the pool News Pool on the HQ cluster. It uses the UNITS command to find out which storage units are available and what their disk numbers are. Then it uses the POOL command to add the storage unit to the pool. Because disk numbers can be different on different nodes, it specifies the name of a node, not the cluster name, on the UNITS and POOL commands.

```
C:\> SNAPMGR HQ1 UNITS
```

Capacity (MB)	Disk Number	Type
4091	1	Shared

```
C:\> SNAPMGR HQ1 POOL "News Pool" /ADD:1
```

A.3.3 Delete a Pool

Syntax

POOL *pool* /DELETE

- *pool* is the name of the pool that you want to delete.

Description

This command deletes a pool, which frees its storage units and brings them online, so that you can access them directly. You can use them to create a new pool, or you can partition and format them, and use them as ordinary disks.

In a cluster, the command fails if the pool cluster resource is offline.

You cannot delete a pool if it contains any virtual disks or snapshots. You must delete all the virtual disks and snapshots in a pool before you can delete the pool.

In a cluster, the command deletes the pool cluster resource and all the storage unit cluster resources. If the pool group is now empty, it deletes the group.

Examples

1. This example deletes the pool News Pool on the HQ cluster.

```
SnapMgr> MANAGE HQ  
SnapMgr> POOL "News Pool" /DELETE
```

2. This example deletes the pool Test on the standalone computer JudyPC.

```
C:\> SNAPMGR JudyPC POOL Test /DELETE
```

A.3.4 Show Pools

Syntax

POOL
POOL *pool*
POOL *pool* /STATISTICS
POOL *pool* /UNITS

- *pool* is the name of the pool about which you want to show information. If you omit this parameter, the command shows information about all the pools that are currently online on the standalone computer or cluster you are managing.

Description

This command shows either:

- Summary information about all pools.
- Full information about one pool.

In a cluster, it does not show pools that are currently offline.

Summary information

When you omit the parameter, the command shows information about all the pools that are currently online on the standalone computer or cluster you are managing. It shows the following information about each pool:

- Name
- Capacity, rounded up to the nearest megabyte
- Free space, rounded down to the nearest megabyte
- Name of the cluster node that currently owns the pool resource. On a standalone computer, this is the name of the standalone computer.

Full information

When you supply a parameter, the command shows full information about the specified pool.

By default, if you do not use the `/STATISTICS` or `/UNITS` switch, the command shows the following additional information about the pool:

- Segment size. A segment is the smallest unit that is copied during a copy-out operation.
- Creation date.
- Last modified; this is when a storage unit was last added to it.
- Version number of its on-disk structures.
- Summary information about all of its virtual disks and snapshots. For more information, see the description of the `VIRTUALDISK` command (Page A-33) or the `SNAPSHOT` command (Page A-21).

Showing I/O statistics

The `/STATISTICS` switch shows information about I/Os to the pool. It shows the following counters, which are zeroed when the computer restarts, or, in a cluster, when the pool fails over within the cluster.

Counter	Description
Copy-outs	The number of segments that were copied to preserve data for snapshots.
Read requests	The number of read I/O requests received by the pool software.
Write requests	The number of write I/O requests received by the pool software.
Disk reads	The number of read I/Os issued to disk by the pool software. This includes reads caused by copy-out operations and split reads.
Disk writes	The number of write I/Os issued to disk by the pool software. This includes writes caused by copy-out operations and split writes.
Split reads	The number of read I/O requests that had to be split into two or more I/Os to disk because the read crossed a segment boundary, and the next segment was not contiguous with the current segment.
Split writes	The number of write I/O requests that had to be split into two or more I/Os to disk because the write crossed a segment boundary, and the next segment was not contiguous with the current segment.

Showing the storage units

The `/UNITS` switch shows the following information about each storage unit in the pool:

- Capacity
- Disk number. In a cluster, this is the disk number on the node that currently owns the pool resource. The disk number may change when you restart the computer or when the pool resource fails over within the cluster.

■ Type:

Shared: a disk in a cluster that is on a shared storage bus and can be seen by every node in the cluster.

Local: a physical disk that is attached to the computer.

Examples

1. This example shows summary information about all the pools that are currently online on the HQ cluster.

```
SnapMgr> MANAGE HQ
SnapMgr> POOL
```

Name	Capacity (MB)	Free Space (MB)	Owner
RAID Pool	17288	8232	HQ1
News Pool	6074	3073	HQ2

2. This example shows information about one pool.

```
SnapMgr> POOL "RAID Pool"
```

```
Pool:          RAID Pool
Capacity:      17288 MB
Free space:    8232 MB
Segment size: 32 KB
Created:       11/15/98 08:38 AM
Modified:      1/29/98 08:45 AM
Version:       1.0
Owner node:    HQ1
```

Name	Type	Capacity (MB)	Delspace (MB)	Drive	Family
Accounts Data	Virtual	4500	0	Z:	1
Engineering Data	Virtual	3000	3000	Y:	2
Accounts Mon	Snapshot	4500	573	W:	1
Accounts Tue	Snapshot	4500	2729	T:	1

3. This example shows I/O statistics for a pool.

```
C:\> SNAPMGR HQ POOL "RAID Pool" /STATISTICS
```

```
Pool:          RAID Pool
Copy-outs:     12345
Read requests: 827652
Write requests: 123456
Disk reads:    73654
Disk writes:   3569
Split reads:   34
Split writes:  11
```

4. This example shows which storage units are in a pool on the HQ cluster.

```
SnapMgr> MANAGE HQ
SnapMgr> POOL "RAID Pool" /UNITS

Capacity (MB)   Disk Number   Type
-----
8678             3             Shared
8678             4             Shared
```

A.4 SNAPSHOT

Summary

Use SNAPSHOT commands to manage snapshots. You can:

- Create a snapshot:
SNAPSHOT *snapshot* /PARENT:*parent*
- Map a drive letter to a snapshot:
SNAPSHOT *snapshot* /MAP:*drive*
- Unmap a drive letter from a snapshot:
SNAPSHOT *snapshot* /UNMAP
- Delete a snapshot:
SNAPSHOT *snapshot* /DELETE [/UNMAP]
- Show information about snapshots:
SNAPSHOT
SNAPSHOT *snapshot*

A.4.1 Create a Snapshot

Syntax

SNAPSHOT *snapshot* /PARENT:*parent*

- *snapshot* is the name that you want to give to the new snapshot. It can be up to 23 characters. If you want to use spaces in the name, enclose it in quotation marks.

You must choose a name that is different from that of any other pool, virtual disk, or snapshot on the computer or cluster you are managing.

If you want to serve the snapshot as a network disk, choose a name that is different from that of any cluster resource, otherwise when you subsequently serve it, you will not be able to make its network disk name the same as its snapshot name.

Note that you cannot change the name of the snapshot later.

- *parent* is the name of the virtual disk or snapshot of which you want to create a snapshot.

Description

This command creates a snapshot of a disk in a pool, called the parent disk. The parent disk can be either a virtual disk or another snapshot. It need not have a drive letter mapped to it.

In a cluster, the command fails if the pool cluster resource is offline.

The command creates a new disk, called a snapshot, that looks like an exact copy of the whole of the parent disk at an instant in time. It has the same capacity and volume label, and it contains the same data.

The snapshot is in the same family as its parent disk. You can have up to 12 snapshots in a family.

To begin with, the snapshot does not take up any disk space. It only starts to take up disk space when you make changes to the data stored on either itself or its parent.

The command flushes the local system cache. Any data in the cache that has not yet been written out to the parent disk is flushed to disk before the snapshot is created.

If the parent disk is being served as a network disk, it does not have any data in the local cache, but may have data in the cache on the remote computer. To flush the remote cache, disconnect the remote computer from the parent disk before you create the snapshot.

The new snapshot is initially offline. It remains offline until you map a drive letter to it. Use the SNAPSHOT command (Page A-21) to map a drive letter to the snapshot, not Disk Administrator.

Examples

1. This example creates a snapshot of the virtual disk Accounts Data on the HQ cluster. The new snapshot is called Accounts Mon.

```
SnapMgr> MANAGE HQ  
SnapMgr> SNAPSHOT "Accounts Mon" /PARENT:"Accounts Data"
```

2. This example creates a snapshot of the virtual disk Scratch on the HQ cluster. The new snapshot is called ScratchSnap.

```
C:\> SNAPMGR HQ SNAPSHOT ScratchSnap /PARENT:Scratch
```

A.4.2 Map a Drive Letter to a Snapshot

Syntax

SNAPSHOT *snapshot* /MAP:*drive*

- *snapshot* is the name of the snapshot to which you want to map a drive letter.
- *drive* is the drive letter you want to map to it.

Description

This command maps a drive letter to a snapshot. The command fails if the pool that the snapshot is in is full.

In a cluster, the command fails if the pool cluster resource is offline.

The drive letter is persistent; the next time the computer restarts, the same drive letter is automatically mapped to the snapshot, provided that the drive letter is available.

If possible, select a letter near the end of the alphabet. During system startup, drive letters are automatically allocated from the start of the alphabet. When the pool software starts up, if the drive letter has already been allocated to another disk that came online earlier in the startup sequence, no drive letter is mapped to the snapshot. The pool software starts quite late in the startup sequence, so select a letter near the end of the alphabet to reduce the probability that the drive letter is allocated to another disk.

In a cluster, if the drive letter is not available when the pool fails over to another node, no drive letter is mapped to the snapshot.

Examples

1. This example maps the drive letter **W** to the snapshot **Accounts Mon** on the **HQ** cluster.

```
SnapMgr> MANAGE HQ  
SnapMgr> SNAPSHOT "Accounts Mon" /MAP:W
```

2. This example maps the drive letter **T** to the snapshot **Accounts Tue** on the **HQ** cluster.

```
C:\> SNAPMGR HQ SNAPSHOT "Accounts Tue" /MAP:T
```

A.4.3 Unmap the Drive Letter from a Snapshot

Syntax

SNAPSHOT *snapshot* /UNMAP

- *snapshot* is the name of the snapshot from which you want to unmap the drive letter.

Description

This command unmaps the drive letter from a snapshot, and takes the snapshot offline. The command fails if any files on the snapshot are open.

In a cluster, the command fails if the pool cluster resource is offline.

In a cluster, the command unmaps the drive letter throughout the cluster; when the pool fails over to another node, the snapshot still has no drive letter.

Examples

1. This example unmaps the drive letter from the snapshot Accounts Mon on the HQ cluster.

```
SnapMgr> MANAGE HQ  
SnapMgr> SNAPSHOT "Accounts Mon" /UNMAP
```

2. This example unmaps the drive letter from the snapshot Accounts Tue on the HQ cluster.

```
C:\> SNAPMGR HQ SNAPSHOT "Accounts Tue" /UNMAP
```

A.4.4 Delete a Snapshot

Syntax

SNAPSHOT *snapshot* /DELETE [/UNMAP]

- *snapshot* is the name of the snapshot that you want to delete.

Description

This command deletes a snapshot from a pool. If you use the /UNMAP switch, it unmaps the drive letter from the snapshot and then deletes it.

You cannot delete a snapshot that has a drive letter mapped to it. If you do not use the /UNMAP switch, the command fails if a drive letter is currently mapped to the snapshot.

Deleting the snapshot frees up the disk space that the snapshot is using that is not shared with any other disks in its family. The amount of free space in the pool increases by the snapshot's Delspace.

The Delspace of other disks in its family may increase:

- If it shared disk space with another snapshot, the Delspace of the other snapshot rises by the amount of the shared disk space.
- If it was the only snapshot in its family, the Delspace of the family's virtual disk changes from 0 to the capacity of the virtual disk. Now you can delete the virtual disk, so its Delspace becomes the same as its capacity.

Deleting the snapshot destroys all the data stored on the snapshot, but doesn't affect the data stored on its parent disk or on any other disks in its family.

The command fails if the snapshot is being served as a network disk.

In a cluster, the command fails if the pool cluster resource is offline for the pool that contains the snapshot.

If you use the /UNMAP switch, the command fails if any files on the snapshot are open.

Examples

1. This example unmaps the drive letter from the snapshot Accounts Mon on the HQ cluster, and then deletes the snapshot.

```
SnapMgr> MANAGE HQ  
SnapMgr> SNAPSHOT "Accounts Mon" /DELETE /UNMAP
```

2. This example unmaps the drive letter from the snapshot Accounts Tue on the HQ cluster, and then deletes the snapshot.

```
C:\> SNAPMGR HQ SNAPSHOT "Accounts Tue" /UNMAP  
C:\> SNAPMGR HQ SNAPSHOT "Accounts Tue" /DELETE
```

A.4.5 Show Snapshots

Syntax

SNAPSHOT

SNAPSHOT *snapshot*

- *snapshot* is the name of the snapshot about which you want to show information. If you omit this parameter, the command shows information about all the snapshots on the computer or cluster you are managing. In a cluster, it shows information only about snapshots that are in pools that are currently online.

Description

This command shows either:

- Summary information about all snapshots.
- Full information about one snapshot.

In a cluster, it shows information only about snapshots that are in pools that are currently online.

Summary information

When you omit the parameter, the command shows information about all the snapshots on the computer or cluster you are managing.

It shows the following information about each snapshot:

- Name
- Name of the pool that it is in
- Drive letter is currently mapped to it.
- Family it is in.
- Name of the cluster node that currently owns its pool. On a standalone computer, this is the name of the standalone computer.

Full information

When you supply a parameter, the command shows full information about the specified snapshot. It shows the following additional information:

- **Capacity.** This is how big the file system and every other piece of software on the system thinks the snapshot is. For example, it is the size shown by Windows NT Explorer.

It is the maximum pool space the snapshot could consume, if you modified the entire contents of either the snapshot or its parent disk.

- **Delspace.** This is the amount of space you would free up in the pool if you deleted it. It is the amount of pool space used exclusively by this snapshot – the space that it is using and that is not shared with any other snapshots in its family.

When there is only one snapshot in a family, the Delspace is the amount of space used by the snapshot.

- **Creation date.**

Examples

1. This example shows summary information about all the snapshots that are in pools that are currently online on the HQ cluster.

```
SnapMgr> MANAGE HQ
SnapMgr> SNAPSHOT
```

Name	Pool	Drive	Family	Owner
Accounts Mon	RAID Pool	W:	1	HQ1
Accounts Tue	RAID Pool	T:	1	HQ1

2. This example shows full information about one snapshot on the HQ cluster.

```
C:\> SNAPMGR HQ SNAPSHOT "Accounts Mon"
```

```
Snapshot:      Accounts Mon
Capacity:      4500 MB
Delspace:      573 MB
Drive letter:  W:
Family:        1
Created:       1/26/99 9:07 PM
Pool:          RAID Pool
```

A.5 UNITS

Syntax

UNITS

Description

This command shows which storage units you could use to create a new pool or add to an existing pool, and what their disk numbers are.

In a cluster, the command fails if any of the nodes in the cluster are down.

The command shows only non-removable disks that are online and contain no partitions. It does not show virtual disks or snapshots. It shows the following information about each storage unit.

Capacity	The capacity of the storage unit.
Disk Number	The disk number of the storage unit. In a cluster, this is the disk number on the node you are managing.
Type	The type of the storage unit: Shared: a disk in a cluster that is on a shared storage bus and can be seen by every node in the cluster. Local: either a physical disk that is attached to the computer, or a network disk.

It shows two types of disk that you must not use in pools:

- Network disks: these are shown as having the type Local. In a cluster, a network disk appears like a local disk on the node that currently owns the SCE Connected Disk resource.
- Disks that are part of volume sets, mirror sets, and stripe sets created using Disk Administrator.

Examples

1. This example shows which storage units you can use to create a pool on the HQ cluster. It specifies a node name, not the cluster name, because disk numbers can be different on different nodes. If you specify the cluster name, the command shows the disk numbers on the node that currently owns the cluster name resource.

```
SnapMgr> MANAGE HQ2  
SnapMgr> UNITS
```

Capacity (MB)	Disk Number	Type
4091	1	Shared
8678	3	Shared
8678	4	Shared
4091	6	Shared
2007	7	Shared

A.6 VIRTUALDISK

Summary

Use VIRTUALDISK commands to manage virtual disks. You can:

- Create a virtual disk:

```
VIRTUALDISK virtualdisk /POOL:pool /CAPACITY:mbytes
```

- Map a drive letter to a virtual disk:

```
VIRTUALDISK virtualdisk /MAP:drive
```

- Partition and format a virtual disk:

```
VIRTUALDISK virtualdisk /FORMAT
```

- Unmap a drive letter from a virtual disk:

```
VIRTUALDISK virtualdisk /UNMAP
```

- Delete a virtual disk:

```
VIRTUALDISK virtualdisk /DELETE [/UNMAP]
```

- Show information about virtual disks:

```
VIRTUALDISK  
VIRTUALDISK virtualdisk
```

A.6.1 Create a Virtual Disk

Syntax

`VIRTUALDISK virtualdisk /POOL:pool /CAPACITY:mbytes`

- *virtualdisk* is the name that you want to give to the new virtual disk. It can be up to 23 characters. If you want to use spaces in the name, enclose it in quotation marks.

You must choose a name that is different from that of any other pool, virtual disk, or snapshot on the computer or cluster you are managing.

If you want to serve the virtual disk as a network disk, choose a name that is different from that of any cluster resource, otherwise when you serve it, you won't be able to make its network disk name the same as its virtual disk name.

Note that you cannot change the name of the virtual disk later

- *pool* is the name of the pool that you want to create the new virtual disk in.
- *mbytes* is the capacity of the new virtual disk. The minimum is 10 MB, and the maximum is equal to the free space in the pool.

Description

This command creates a new virtual disk of the specified capacity in the specified pool.

In a cluster, the command fails if the pool cluster resource is offline.

The new virtual disk is the first disk in a new family. The SWVR software automatically allocates a number from 0 to 7 to the new family. You can create up to eight virtual disks in a pool.

The virtual disk is initially offline. It remains offline until you map a drive letter to it. Use the VIRTUALDISK command (Page A-33) to map a drive letter to the virtual disk, not Disk Administrator.

Examples

1. This example creates a 300 MB virtual disk called TempDisk in the RAID Pool on the HQ cluster.

```
SnapMgr> MANAGE HQ  
SnapMgr> VIRTUAL TempDisk /POOL:"RAID Pool" /CAPACITY:300
```

2. This example creates a 1000 MB virtual disk called Usenet in the News Pool on the local cluster.

```
C:\> SNAPMGR VIRTUAL Usenet /POOL:"News Pool" /CAPACITY:1000
```

A.6.2 Map a Drive Letter to a Virtual Disk

Syntax

VIRTUALDISK *virtualdisk* /MAP:*drive*

- *virtualdisk* is the name of the virtual disk to which you want to map a drive letter.
- *drive* is the drive letter you want to map to it.

Description

This command maps a drive letter to a virtual disk. The command fails if the pool that the disk is in is full.

In a cluster, the command fails if the pool cluster resource is offline.

The drive letter is persistent; the next time the computer restarts, the same drive letter is automatically mapped to the virtual disk, provided that the drive letter is available.

If possible, select a letter near the end of the alphabet. During system startup, drive letters are automatically allocated from the start of the alphabet. When the pool software starts up, if the drive letter has already been allocated to another disk that came online earlier in the startup sequence, no drive letter is mapped to the virtual disk. The pool software starts quite late in the startup sequence, so select a letter near the end of the alphabet to reduce the probability that the drive letter is allocated to another disk.

In a cluster, if the drive letter is not available when the pool fails over to another node, no drive letter is mapped to the virtual disk.

Examples

1. This example maps the drive letter Z to the virtual disk Accounts Data on the HQ cluster.

```
SnapMgr> MANAGE HQ  
SnapMgr> VIRTUALDISK "Accounts Data" /MAP:Z
```

2. This example maps the drive letter Y to the virtual disk Engineering Data on the HQ cluster.

```
C:\> SNAPMGR HQ VIRTUALDISK "Engineering Data" /MAP:Y
```

A.6.3 Partition and Format a Virtual Disk

Syntax

VIRTUALDISK *virtualdisk* /FORMAT

- *virtualdisk* is the name of the virtual disk that you want to partition and format.

Description

This command creates a single partition on a virtual disk and formats it with the NTFS file system, giving it the same volume label as the name of the virtual disk.

You must have previously mapped a drive letter to the virtual disk.

In a cluster, the command fails if the pool cluster resource is offline.

Examples

1. This example partitions and formats the virtual disk Accounts Data on the HQ cluster.

```
SnapMgr> MANAGE HQ  
SnapMgr> VIRTUALDISK "Accounts Data" /FORMAT
```

2. This example partitions and formats the virtual disk Engineering Data on the HQ cluster.

```
C:\> SNAPMGR HQ VIRTUALDISK "Engineering Data" /FORMAT
```

A.6.3 Unmap the Drive Letter from a Virtual Disk

Syntax

VIRTUALDISK *virtualdisk* /UNMAP

- *virtualdisk* is the name of the virtual disk from which you want to unmap the drive letter.

Description

This command unmaps the drive letter from a virtual disk and takes the virtual disk offline. The command fails if any files on the virtual disk are open.

In a cluster, the command fails if the pool cluster resource is offline.

In a cluster, the command unmaps the drive letter throughout the cluster; when the pool fails over to another node, the virtual disk still has no drive letter.

Examples

1. This example unmaps the drive letter from the virtual disk Accounts Data on the HQ cluster.

```
SnapMgr> MANAGE HQ  
SnapMgr> VIRTUALDISK "Accounts Data" /UNMAP
```

2. This example unmaps the drive letter from the virtual disk Engineering Data on the HQ cluster.

```
C:\> SNAPMGR HQ VIRTUALDISK "Engineering Data" /UNMAP
```

A.6.4 Delete a Virtual Disk

Syntax

VIRTUALDISK *virtualdisk* /DELETE [/UNMAP]

- *virtualdisk* is the name of the virtual disk that you want to delete.

Description

This command deletes a virtual disk from a pool. If you use the /UNMAP switch, it unmaps the drive letter from the virtual disk and then deletes it.

You cannot delete a virtual disk that has a drive letter mapped to it. If you do not use the /UNMAP switch, the command fails if a drive letter is currently mapped to the virtual disk.

Deleting the virtual disk destroys all the data on the disk. It frees up the disk space used by the virtual disk. The free space in the pool increases by the capacity of the virtual disk.

You cannot delete a virtual disk if either:

- There are any snapshots in its family.
- It is being served as a network disk.

In a cluster, the command fails if the pool cluster resource is offline for the pool that the virtual disk is in.

If you use the /UNMAP switch, the command fails if any files on the virtual disk are open.

Examples

1. This example unmaps the drive letter from the virtual disk TempDisk on the HQ cluster, and then deletes the virtual disk.

```
SnapMgr> MANAGE HQ  
SnapMgr> VIRTUALDISK TempDisk /DELETE /UNMAP
```

2. This example unmaps the drive letter from the virtual disk Scratch on the local cluster, and then deletes the virtual disk.

```
C:\> SNAPMGR VIRTUALDISK Scratch /UNMAP  
C:\> SNAPMGR VIRTUALDISK Scratch /DELETE
```

A.6.5 Show Virtual Disks

Syntax

VIRTUALDISK

VIRTUALDISK *virtualdisk*

- *virtualdisk* is the name of the virtual disk about which you want to show information. If you omit this parameter, the command shows information about all the virtual disks on the computer or cluster you are managing. In a cluster, it only shows information about virtual disks that are in pools that are currently online.

Description

This command shows either:

- Summary information about all virtual disks.
- Full information about one virtual disk.

In a cluster, it shows information only about virtual disks that are in pools that are currently online.

Summary information

When you omit the parameter, the command shows information about all the virtual disks on the computer or cluster you are managing.

It shows the following information about each virtual disk:

- Name
- Name of the pool that it is in
- Drive letter is currently mapped to it
- Family that it is in
- Name of the cluster node that currently owns its pool. On a standalone computer, this is the name of the standalone computer.

Full information

When you supply a parameter, the command shows full information about the specified virtual disk. It shows the following additional information:

- Capacity
- Delspace. This is the amount of space you would free in the pool if you deleted it. It is the same as the capacity of the virtual disk, unless there are any snapshots in its family, in which case it is 0. When a virtual disk has a snapshot in its family, you cannot delete the virtual disk, and so its Delspace is 0.
- Creation date

Examples

1. This example shows summary information about all the virtual disks that are in pools that are currently online in the HQ cluster.

```
SnapMgr> MANAGE HQ  
SnapMgr> VIRTUALDISK
```

Name	Pool	Drive	Family	Owner
Accounts Data	RAID Pool	Z:	1	HQ1
Engineering Data	RAID Pool	Y:	2	HQ1
Usenet News	News Pool	X:	0	HQ2

2. This example shows full information about one virtual disk on the HQ cluster.

```
C:\> SNAPMGR HQ VIRTUALDISK "Accounts Data"
```

```
Virtual disk: Accounts Data  
Capacity: 4500 MB  
Delspace: 0 MB  
Drive letter: Z:  
Family: 1  
Created: 1/26/99 9:07 PM  
Pool: RAID Pool
```


B NDMgr Commands

This appendix describes the commands that you can use to manage network disks.

The following table gives a summary of the commands. For information on how to use them, see “Using the Command Line Interface” on Page 5-6.

Command	Description
DRIVES	Shows which drive letters are available for you to map to virtual disks, snapshots, and the network disks to which you connect.
MANAGE	Controls which computer or cluster you manage (only available at the SnapMgr or NDMgr prompt).
REGISTER	Manages the list of client computers that are allowed to connect to the network disks that you serve.
SERVE	Manages the network disks that you serve.
USE	Manages the network disks to which you connect.

B.1 DRIVES

Syntax

DRIVES [/NODE:*nodename*]

- *nodename* is the name of the cluster node whose available drive letters you want to show.

Description

Use the DRIVES command to show which drive letters are available for you to map to virtual disks, snapshots, and the network disks to which you connect.

In a cluster, the command shows drive letters that are available only on every node in the cluster.

If one of the nodes in the cluster is down, use the /NODE switch to find out which drive letters are available on a particular node. The DRIVES command fails if one of the nodes is down and you do not use the /NODE switch.

Examples

1. This example shows which drive letters are available on every node in the HQ cluster.

```
NDMgr> MANAGE HQ  
NDMgr> DRIVES
```

```
Available drive letters: DGHMNPQ
```

2. This example shows which drive letters are available on the standalone computer Backup.

```
C:\> NDMGR Backup DRIVES
```

```
Available drive letters: DEFGHMNPQRSTVW
```

3. This example shows which drive letters are available on the local cluster node Accounts1.

```
NDMgr> MANAGE /LOCAL  
NDMgr> DRIVES /NODE:Accounts1
```

```
Available drive letters: GHKMNPQRST
```

B.2 MANAGE

Summary

Use MANAGE commands to control which computer or cluster you manage. You can:

- Manage a remote computer or cluster:
MANAGE *computername*
- Manage the local computer or cluster:
MANAGE /LOCAL
- Show which computer or cluster you are managing:
MANAGE

You can use MANAGE commands only at the SnapMgr or NDMgr prompt. You cannot use them at the Windows NT command prompt.

If you are entering commands at the Windows NT command prompt, each command automatically manages the local computer or cluster unless you specify the name of a remote computer or cluster after the SNAPMGR or NDMGR command prefix.

For example, these commands show which drive letters are available, first on the Accounts cluster, and then on the local cluster:

```
C:\> NDMGR Accounts DRIVES  
C:\> NDMGR DRIVES
```

B.2.1 Manage a Remote Computer

Syntax

MANAGE *computername*

- *computername* is the name of the standalone computer or cluster that you want to manage.

Description

You can use this command only at the SnapMgr or NDMgr prompt.

By default, you manage the local standalone computer or cluster. Use this command to manage a remote computer or cluster. Subsequent commands that you enter at the SnapMgr or NDMgr prompt will manage the specified computer or cluster.

Examples

1. This example manages the standalone computer called Usenet. Subsequent commands that you enter at the NDMgr prompt will manage Usenet.

```
NDMgr> MANAGE Usenet
```

B.2.2 Manage the Local Computer

Syntax

MANAGE /LOCAL

Description

You can use this command only at the SnapMgr or NDMgr prompt.

If you have previously issued a MANAGE command to manage a remote computer or cluster, use this command to revert to managing the local computer or cluster. Subsequent commands that you enter at the SnapMgr or NDMgr prompt will manage the local standalone computer or cluster.

Examples

1. This example shows which drive letters are available, first on the Accounts cluster, and then on the local cluster.

```
NDMgr> MANAGE Accounts
NDMgr> DRIVES
NDMgr> MANAGE /LOCAL
NDMgr> DRIVES
```

B.2.3 Show the Managed Computer

Syntax

MANAGE

Description

You can use this command only at the SnapMgr or NDMgr prompt.

This command shows which standalone computer or cluster you are currently managing. This is the computer or cluster that the subsequent commands that you enter at the SnapMgr or NDMgr prompt will manage.

Examples

1. In this example, you are managing the standalone computer called Usenet. Subsequent commands issued at the NDMgr prompt will manage Usenet.

```
NDMgr> MANAGE
Managing node USENET
```

2. In this example, you are managing the Accounts cluster. Subsequent commands issued at the NDMgr prompt will manage the Accounts cluster.

```
NDMgr> MANAGE
Managing cluster ACCOUNTS
Cluster members: ACCOUNTS1      ACCOUNTS2
```

B.3 REGISTER

Summary

Use REGISTER commands to manage the list of client computers that are allowed to connect to the network disks served by the computer or cluster you are managing. You can:

- Add a registered client:
REGISTER *clientname*
- Remove a registered client:
REGISTER *clientname* /DELETE
- Show registered clients:
REGISTER

B.3.1 Add a Registered Client

Syntax

REGISTER *clientname*

- *clientname* is the name of the computer that you want to add to the list of registered clients.

Description

This command adds an entry to the list of registered clients on the standalone computer or cluster you are managing. The list specifies which computers are allowed to connect to the network disks served by that computer or cluster.

By default, the list is empty, which means that all computers can connect to the network disks served by the computer or cluster you are managing.

If the list is **not** empty, only computers on the list can connect to its network disks. When a remote computer tries to connect to one of these network disks, the request is refused if the name of the computer is not on the list.

Note that **any** computer on the list can connect to **any** of its network disks. You cannot specify which computers can connect to which network disks.

If you want a cluster to be able to connect to the network disks served by the computer or cluster you are managing:

- Do not add the name of the cluster, because cluster names are ignored in the list of registered clients.
- Add the name of each computer in the cluster. Add one entry for each node in the cluster.

Note that the command does not check if the specified computer exists.

In a cluster, the command updates the lists on every node in the cluster. It adds the same entry to the list on every node that is currently up.

Examples

1. This example adds the computer Accounts1 to the list of registered clients on the HQ cluster.

```
NDMgr> MANAGE HQ  
NDMgr> REGISTER Accounts1
```

2. This example adds the computer Accounts2 to the list of registered clients on the HQ cluster.

```
C:\> NDMGR HQ REGISTER Accounts2
```

B.3.2 Remove a Registered Client

Syntax

REGISTER *clientname* /DELETE

- *clientname* is the name that you want to remove from the list of registered clients.

Description

This command removes an entry from the list of registered clients on the standalone computer or cluster you are managing. The list specifies which computers are allowed to connect to the network disks served by that computer or cluster.

If a computer with the specified name is currently connected to a network disk served by the computer or cluster you are managing, that computer continues using the disk until it either disconnects from it or loses the connection.

Note that if you remove all the entries from the list, any computer can now connect to the network disks served by the computer or cluster you are managing.

In a cluster, the command updates the list on every node in the cluster. It removes the entry from the list on every node that is currently up.

Examples

1. This example removes the computer Accounts1 from the list of registered clients on the HQ cluster.

```
NDMgr> MANAGE HQ  
NDMgr> REGISTER Accounts1 /DELETE
```

2. This example removes the computer Accounts2 from the list of registered clients on the HQ cluster.

```
C:\> NDMGR HQ REGISTER Accounts2 /DELETE
```

B.3.3 Show Registered Clients

Syntax

REGISTER

Description

This command shows the list of registered clients on the standalone computer or cluster you are managing. The list specifies which computers are allowed to connect to the network disks served by that computer or cluster.

Cluster names in the list are ignored.

Examples

1. This example shows the list of registered clients on the local cluster. The list is empty, so any computer can connect to the network disks served by the local cluster.

```
NDMgr> MANAGE /LOCAL
NDMgr> REGISTER
```

```
Clients registered to connect to network disks served by this
computer:
```

```
(all)
```

2. This example shows the list of registered clients on the HQ cluster.

```
NDMgr> MANAGE HQ
NDMgr> REGISTER
```

```
Clients registered to connect to network disks served by this
cluster:
```

```
ACCOUNTS1
ACCOUNTS2
BACKUP
ENGINEERING1
ENGINEERING2
USENET
```

3. In this example, the lists are different on the nodes in the Central cluster.

```
C:\> NDMGR Central REGISTER
```

```
Warning: The list of registered clients is different on one or more members of this cluster. Each member's list will be displayed.
```

```
Clients registered to connect to network disks served by member CENTRAL1:
```

```
(all)
```

```
Clients registered to connect to network disks served by member CENTRAL2:
```

```
ACCOUNTS1  
ACCOUNTS2
```

B.4 SERVE

Summary

Use SERVE commands to manage the network disks that your computer or cluster serves over the network. You can:

- Serve a disk:
SERVE *netdiskname drive*
- Stop serving a disk:
SERVE *netdiskname /DELETE*
- Show served disks:
SERVE
SERVE *netdiskname*

B.4.1 Serve a Disk

Syntax

SERVE *netdiskname drive*

- *netdiskname* is the name that you want to give to the network disk. It can be up to 23 characters. If you want to use spaces in the name, enclose it in quotation marks. The name you supply is automatically converted to uppercase.

You must choose a name that is different from that of any other network disk that is being served by the computer or cluster. In a cluster, the name must also be different from that of any existing cluster group or cluster resource.

We recommend that you make the name the same as the disk's volume label, so that you can easily recognize it when you use Windows NT Explorer on the client computer. If you are serving a virtual disk or snapshot, make sure that its virtual disk name or snapshot name is also the same as its label.

- *drive* is the drive letter currently mapped to the disk you want to serve.

Description

This command serves a disk on the standalone computer or cluster you are managing. The command marks the disk as offline on that computer or cluster, so that its users can no longer access the disk. This prevents unsynchronized access between these users and users on the client system that subsequently connects to the disk.

- ▲ **Note:** Make sure that all the files on the disk are closed before you serve it. The disk is marked as offline when you serve it, which means that users who were accessing the files could lose data.

The disk you are serving should have only one partition, which should be formatted with the NTFS file system.

The standalone computer or cluster you are managing automatically serves the disk whenever it restarts, until you explicitly stop serving the disk, using either the `SERVE` command or the Network Disk Manager or Snapshot Manager `snap-in`.

If it fails to serve the disk during startup, for example, because you physically removed the disk, it does not try to serve the disk again until the next time it restarts.

Before you serve a standard disk (not a virtual disk or snapshot), follow these steps to make sure that the drive letter mapped to the disk is persistent:

1. Choose which drive letter you want to map to the disk.
In a cluster, use the `NDMgr` or `SnapMgr DRIVES` command (page B-2) to find out which drive letters are available throughout the cluster. Choose one of those drive letters.
2. Use Disk Administrator to map the drive letter to the disk and make it persistent.
Select the disk and deassign its current drive letter, even if it already has the same drive letter as the one you chose in Step 1. Now assign the drive letter you chose in Step 1 to the disk. The drive letter is now persistent on the local computer.
3. In a cluster, repeat Step 2 on every node that is attached to the disk.
The disk now has a persistent drive letter throughout the cluster.

Cluster resources

Before you serve a disk in a cluster, make sure that it is not configured as a cluster resource of type `Physical Disk`. Use Cluster Administrator to check whether the disk you want to serve is configured as a cluster resource of type `Physical Disk`, and to delete the resource if necessary.

When you serve a disk in a cluster, the command automatically creates a cluster resource of type `SCE Served Disk` and brings the resource online. The resource has the same name as the network disk. The actions taken by the command depend on whether you are serving a virtual disk or snapshot, or a standard disk:

■ **Serving a virtual disk or snapshot**

The command creates the SCE Served Disk resource in the pool's cluster group, and makes it dependent on the pool resource.

- ▲ **Note:** Do not use Cluster Administrator to change the dependency of the served disk resource on the pool resource.

■ **Serving other disks**

The command creates a new cluster group for the SCE Served Disk resource. The group has the same name as the network disk.

If the disk you served cannot be seen by all the nodes in the cluster, use Cluster Administrator to restrict its possible owners to the nodes that can see it. For example, if you served a local disk, restrict the possible owners to the one computer that is attached to the disk.

- ▲ **Note:** If the SCE Served Disk resource is allowed to fail over to a node that cannot see the disk, another disk that happens to have the same drive letter on the target node could be served in its place.

Examples

1. This example serves drive Z on the HQ cluster and gives it the network disk name Accounts Data. Drive Z is the virtual disk Accounts Data; notice that we give the network disk the same name as the virtual disk.

```
NDMgr> MANAGE HQ  
NDMgr> SERVE "Accounts Data" Z:
```

2. This example serves drive W on the HQ cluster and gives it the network disk name Accounts Mon. Drive W is the snapshot Accounts Mon; notice that we give the network disk the same name as the snapshot.

```
C:\> NDMGR HQ SERVE "Accounts Mon" W:
```

B.4.2 Stop Serving a Disk

Syntax

SERVE *netdiskname* /DELETE

- *netdiskname* is the name of the network disk that you want to stop serving.

Description

This command stops the standalone computer or cluster you are managing from serving a disk. The command brings the disk back online on that computer or cluster, so that its users can access it again, for example, by double-clicking its icon in Windows NT Explorer.

Disconnect the client computer from the network disk before you stop serving it.

You cannot stop serving a disk if a computer is currently connected to it. But if a cluster is connected to it, you could accidentally stop serving it if the SCE Connected Disk resource is currently failing over on the client cluster. Be careful not to inadvertently stop serving a disk that is failing over on the client cluster.

In a cluster, the command deletes the SCE Served Disk resource that is associated with the network disk. It also deletes the group that the resource was in if that group is now empty.

In a cluster, the command fails if any node in the cluster is down.

Examples

1. This example stops serving the network disk Accounts Mon on the HQ cluster:

```
NDMgr> MANAGE HQ
NDMgr> SERVE "Accounts Mon" /DELETE
```

2. This example stops serving the network disk Accounts Data on the HQ cluster:

```
C:\> NDMGR HQ SERVE "Accounts Data" /DELETE
```

B.4.3 Show Served Disks

Syntax

SERVE
SERVE *netdiskname*

- *netdiskname* is the name of the network disk about which you want to show information. If you omit this parameter, the command shows information about all network disks that are being served by the standalone computer or cluster you are managing. In a cluster, it shows only network disks whose SCE Served Disk resources are currently online.

Description

This command shows either:

- Summary information about all served disks.
- Full information about one served disk.

In a cluster, it shows information only about network disks whose SCE Served Disk resources are currently online.

Summary information

When you omit the parameter, the command shows information about all the network disks that are being served by the standalone computer or cluster you are managing.

It shows the following information about each network disk:

- Network disk name

- **Disk number.** In a cluster, this is the disk number on the node that currently owns the SCE Served Disk resource associated with the disk. The disk number may change when you restart the computer, or when the SCE Served Disk resource fails over within the cluster. The value **Unknown** is shown if the computer or cluster failed to serve the network disk at startup because the disk had been removed. In this case, it will not try to serve the disk again until the next time it restarts.
- **Name of the client computer** that is currently connected to the network disk. If the client is a cluster, this is the name of the node that currently owns the SCE Connected Disk resource, not the name of the cluster.
It is blank if no computer is connected. It is also blank if either:
 - A standalone computer or cluster was connected but lost the connection because the service was temporarily unavailable, and is now trying to reconnect.
 - The client is a cluster and the SCE Connected Disk resource is currently failing over in the client cluster.

Full information

When you supply a parameter, the command shows full information about the specified network disk.

In a cluster, it shows the name of the node that currently owns the SCE Served Disk resource.

It also shows the following I/O statistics.

Counter	Description
Read operations	The number of read I/Os issued to the served disk.
Data read	The amount of data read from the served disk.
Write operations	The number of write I/Os issued to the served disk.
Data written	The amount of data written to the served disk.

The counters are zeroed when a standalone computer restarts.

In a cluster, they are zeroed when the disk comes online on a node (for example, at cluster failover) if the disk is a served virtual disk or snapshot. If the disk is **not** a virtual disk or snapshot, the counters are

zeroed when the cluster restarts, and they measure cumulative information for that particular node since cluster startup.

Examples

1. This example shows summary information about all the network disks that are currently online and being served by the HQ cluster.

```
NDMGR> MANAGE HQ
NDMGR> SERVE
```

This cluster is configured to serve 4 network disks:

Network Disk	Disk Number	Client
ACCOUNTS DATA	13	ACCOUNTS1
ACCOUNTS MON	11	BACKUP
ENGINEERING DATA	14	ENGINEERING2
USENET NEWS	12	USENET

2. This example shows full information about the network disk Accounts Data that is being served by the HQ cluster.

```
C:\> NDMGR HQ SERVE "Accounts Data"
```

```
Network disk:      \\HQ\ACCOUNTS DATA
Client:           ACCOUNTS1
Disk number:      13
Cluster resource: Yes
Owner:           HQ1
Read operations:  18091
Data read:        103.37 MB
Write operations: 7394
Data written:     530.90 MB
```

B.5 USE

Summary

Use USE commands to manage the network disks that are connected to your computer or cluster. You can:

- Connect to a network disk:

- USE *drive* \\servername\netdiskname

- USE * \\servername\netdiskname

- Disconnect from a network disk:

- USE \\servername\netdiskname /DELETE

- USE *drive* /DELETE

- Show information about connected network disks:

- USE

- USE \\servername\netdiskname

- USE *drive*

B.5.1 Connect to a Network Disk

Syntax

USE *drive* \\servername\netdiskname

USE * \\servername\netdiskname

- *drive* is the drive letter that you want to temporarily map to the disk.
- *servername* is the name of the computer or cluster that is serving the network disk to which you want to connect. If it is being served by a cluster, you must specify the name of the cluster; you cannot specify the name of a node in the cluster.
- *netdiskname* is the name of the network disk to which you want to connect.

Description

This command connects the standalone computer or cluster you are managing to a network disk. It creates a new disk on that computer or cluster, which behaves like a locally attached disk. Applications running on it are not aware that the disk is actually located elsewhere on the network.

The computer or cluster automatically tries to reconnect to the network disk whenever it restarts, until you explicitly disconnect it, using either the USE command (Page B-23), or the Network Disk Manager or Snapshot Manager snap-in.

If it fails to connect during startup, for example because the network link is down, it doesn't try to connect again until the next time it restarts. In this case:

- On a standalone computer, the snap-in and USE command (Page B-23) show the network disk status as Failed.
- In a cluster, the snap-in and USE command (Page B-23) do not show the disk. Cluster Administrator shows the SCE Connected Disk resource as either failed or offline.

The command fails if another computer is already connected to the network disk. Only one computer at a time can connect to a network disk.

Mapping a drive letter

The command automatically maps a drive letter to the network disk.

If the first parameter of the command is an asterisk (*), it maps a drive letter as follows:

- If this is the first time that the standalone computer or cluster you are managing has connected to the network disk, the command temporarily maps the first available drive letter.
- If the standalone computer or cluster you are managing has previously connected to the network disk, **and** you made its drive letter persistent, the same drive letter is mapped to it now, provided that the drive letter is available. If another disk or file share is already using the drive letter, the first available drive letter is temporarily mapped to the network disk.

If you are managing the local standalone computer or cluster, you can use the *drive* parameter to specify which drive letter you would like to temporarily map to the network disk. Use Disk Administrator to make it persistent.

If you are managing a remote standalone computer or cluster, the *drive* parameter is ignored. It behaves like the asterisk (*) parameter.

In a cluster

In a cluster, the command automatically creates a cluster group and a cluster resource in that group. Both the group and the resource have the name `\\servername\netdiskname`, where:

- *servername* is the name of the standalone computer or cluster that is serving the disk
- *netdiskname* is the name of the network disk

The new resource is of type SCE Connected Disk.

- ▲ **Note:** Do not use Cluster Administrator to change the name of the SCE Connected Disk resource.

What happens if the computer loses the connection?

This section looks at what happens when the client computer subsequently loses the connection to the network disk because the service is temporarily unavailable:

- Because the network link failed, or
- Because the server failed (either the computer serving the disk failed, or, if it's being served by a cluster, the served disk resource is failing over from one node in that cluster to another).

What happens depends on whether the network disk is a cluster quorum disk or not.

Network disk is not cluster quorum disk

If the network disk is not being used as a cluster quorum disk, when the computer loses the connection, it logs a warning record in the system event log then automatically tries to reconnect. When it manages to reconnect, it logs an information message. Both of these records have NTDS in the event record's *Source* field.

If the computer does not manage to reconnect immediately, it keeps on trying, because it may have data in its cache that needs to be written to disk. If the computer is configured to use multiple network links for network disk connections, it automatically tries to reconnect using all of those network links.

Until the computer manages to reconnect to the network disk, it stalls all read and write I/O requests from the system cache to the network disk. If the cache becomes full because it cannot issue writes to the network disk, the applications using the network disk hang until the network disk becomes available again.

If the computer fails or you turn its power off before it manages to reconnect, you lose the data in its cache that has not yet been written to disk.

If the computer is a standalone computer and you try to shut it down before it manages to reconnect, the shutdown will hang until the computer has reconnected and written any data in its cache to disk. If the shutdown hangs, you have two choices:

- Fix the problem that made the network disk unavailable, for example, repair the network link. The shutdown will then proceed as normal.

- Turn off the power and lose the data in the cache.

Network disk is cluster quorum disk

If the network disk is being used as the cluster quorum disk, when the node that is currently connected to the network disk loses the connection, the node may be removed from the cluster.

During a cluster IsAlive check, if the status of the network disk is Reconnecting, the node is disconnected from the network disk and removed from the cluster.

The computer then behaves in the normal way following removal from a cluster. It waits for 5 minutes then attempts to rejoin the cluster, and so on.

Examples

1. This example connects the network disk Scratch to the local computer and temporarily maps the first available drive letter to the disk. The disk is being served by the HQ cluster.

```
NDMgr> MANAGE /LOCAL  
NDMgr> USE * \\HQ\Scratch
```

2. This example connects the Accounts cluster to the network disk Accounts Data and temporarily maps the drive letter Y to the disk. The disk is being served by the HQ cluster.

```
C:\> NDMGR Accounts USE Y: \\HQ\“Accounts Data”
```

B.5.2 Disconnect from a Network Disk

Syntax

USE \\servername\netdiskname /DELETE

USE drive /DELETE

- *servername* is the name of the computer or cluster that is serving the network disk from which you want to disconnect. If it is being served by a cluster, you must specify the name of the cluster; you cannot specify the name of a node in the cluster.
- *netdiskname* is the name of the network disk.
- *drive* is the drive letter currently mapped to the network disk.

Description

This command disconnects the computer or cluster you are managing from a network disk.

If the status of the network disk is Connected on that computer or cluster, the local system cache is automatically flushed. All the data that has not yet been written to the disk is flushed to disk before it is disconnected.

If the status is currently Reconnecting (the computer or cluster has lost the connection and trying to reconnect):

- On a standalone computer, the attempt to disconnect will hang, if necessary forever, until the computer reconnects and writes its cached data to disk.
- In a cluster, the software continues trying to reconnect for 30 seconds. If the connection has not been re-established within 30 seconds, the disk is disconnected and data in the cache that had not been written to disk is lost.

In a cluster, the command deletes the SCE Connected Disk resource that is associated with the network disk. It also deletes the group the resource was in, if that group is now empty.

In a cluster, the command fails if there are any cluster resources that depend on the SCE Connected Disk resource.

Examples

1. This example disconnects the local computer from the network disk Scratch. The network disk is being served by the HQ cluster.

```
NDMgr> MANAGE /LOCAL  
NDMgr> USE \\HQ\Scratch /DELETE
```

2. This example disconnects the Backup computer from the network disk Accounts Mon. The network disk is being served by the HQ cluster.

```
C:\> NDMGR Backup USE \\HQ"Accounts Mon" /DELETE
```

B.5.3 Show Connected Network Disks

Syntax

USE

USE *\\servername\netdiskname*

USE *drive*

- *servername* is the name of the computer or cluster that is serving the network disk about which you want to show information. If it is being served by a cluster, you must specify the name of the cluster; you cannot specify the name of a node in the cluster.
- *netdiskname* is the name of the network disk about which you want to show information.
- *drive* is the drive letter that is currently mapped to the network disk about which you want to show information.

Description

This command shows either:

- Summary information about all network disks.
- Full information about one network disk.

In a cluster, it shows information only about network disks whose SCE Connected Disk resources are currently online.

Summary information

When you omit the parameter, the command shows information about all the network disks to which the computer or cluster you are managing is connected.

It shows the following information about each network disk:

- Its network disk name and the name of the computer or cluster that is serving it.

- Its status (see the following table).
- Its disk number. In a cluster, this is the disk number on the node that currently owns the SCE Connected Disk resource associated with the disk. The disk number may change when you restart the computer, or when the connected disk resource fails over in a cluster.
- The drive letter currently mapped to it.

Status	Description
Connecting	In the process of connecting.
Connected	Connected.
Disconnecting	In the process of disconnecting.
Disconnected	Disconnected.
Failed	Failed to connect at last startup, because either the network was down or the server was unavailable at the time; will try to connect again at next startup. You see this status only on a standalone computer.
Reconnecting	<p>Currently trying to reconnect; the computer lost the connection because the service was unavailable (either the server or the network link failed, or the disk failed over within the server cluster).</p> <p>In a cluster, the node that currently owns the connected disk resource keeps trying to reconnect, if necessary forever. If it lost the connection because the network link failed, the failure may not affect every node in the cluster. If another node can still see the system serving the network disk, use Cluster Administrator to move the SCE Connected Disk resource to that node.</p>

Full information

When you supply a parameter, the command shows full information about the specified network disk.

In a cluster, it shows the name of the node that currently owns the SCE Connected Disk resource.

Examples

1. This example shows summary information about all the network disks that the Accounts cluster is configured to connect to:

```
NDMgr> MANAGE Accounts
NDMgr> USE
```

This cluster is configured to connect to 1 network disk:

Network Disk	Status	Disk	Drive
\\HQ\ACCOUNTS DATA	Connected	7	Y:

2. This example shows information about a single network disk:

```
NDMgr> USE Y:
```

```
Network disk:      \\HQ\ACCOUNTS DATA
Disk number:      7
Status:           Connected
Drive:            Y:
Cluster resource: Yes
Owner:            ACCOUNTS2
```

Glossary

availability

Measures the proportion of time that a system (cluster or standalone computer) provides the service expected by its users. A highly available system provides continuous service despite the failure of its component parts.

cluster

Two or more independent computers that are accessed and managed as a single system. *See also* MSCS cluster.

Cluster Administrator

A management tool provided with MSCS clusters that you use to configure and manage clusters.

cluster group

A collection of one or more related cluster resources that are managed as a single entity. A group is the unit of failover in the cluster.

cluster name

A name that identifies a cluster that you specify when you install the cluster software. You can change it later, using Cluster Administrator. It represents one or more computers that are treated as a single entity.

cluster resource

An MSCS entity that corresponds to a Windows NT resource, such as a physical disk, application, file share, or TCP/IP address. The cluster resources maintain the state, failover properties and interdependencies of Windows NT resources in a cluster.

computer name

The name shown on the Identification tab when you click the Network icon on the computer's Control Panel.

copy-out operation

An operation that happens when two disks in the same family share a segment of disk space and you modify the data stored in the segment. The SWVR software automatically copies the segment to a new location in the pool before it updates the segment, so that the original segment is preserved.

Delspace

Each virtual disk and snapshot in a pool has a Delspace value, which is the amount of space you would free up in the pool if you deleted the virtual disk or snapshot:

- **virtual disk:** The Delspace is the same as the capacity of the virtual disk, unless there are any snapshots in its family. When a virtual disk has a snapshot in its family, you cannot delete the virtual disk, and so its Delspace is 0.
- **snapshot:** The Delspace is the amount of space used exclusively by the snapshot – the space that it is using and that is not shared with any other disks in its family.

The Delspace value is rounded down to the nearest MB.

disk number

A number that is allocated to a disk by Windows NT and that is displayed by Disk Administrator. In a cluster, a disk can have different disk numbers on different nodes in the cluster.

extended LAN

A network that is physically distributed over a wide area, like a WAN, but that uses bridges instead of routers to control network traffic.

failover

The process of transferring the ownership of a cluster group or resource from one node in a cluster to another. The resource or group is taken offline on one node, then brought online on the target node.

family

A set of disks in a pool that are derived from one another and that can share disk space.

A family consists of a virtual disk, its snapshots, the snapshots of those snapshots, and so on. You can have up to 8 families in a pool. Each family is allocated a number, from 0 to 7, by the SWVR software.

hub

Network hardware that joins communications lines together in a star configuration.

local disk

A disk that is physically attached to just one node in a cluster.

metadata

Data stored on disk that is used internally by software to manage and keep track of users' data. Metadata is data *about* data.

MIB

Management Information Base. A collection of managed objects in a database.

Microsoft Management Console (MMC)

A framework for running the management tools provided by Microsoft and other vendors. It allows you to manage local and remote computers.

MSCS cluster

A cluster that is created and managed by the Microsoft Cluster Server (MSCS) software. Also sometimes called Wolfpack clusters, after the internal name used when the software was being developed. *See also* cluster.

names

See cluster name, computer name, network disk name, snapshot name, virtual disk name.

network disk

A disk located on one computer or cluster that is served over the network. When another computer or cluster (the client) connects to the disk, it behaves like a disk that is physically attached to the client. Applications

running on the client are not aware that the disk is located elsewhere on the network.

Network Disk Manager

A tool that you use to manage network disks. It is an MMC snap-in. To run it, double-click the appropriate MMC console file.

A default console file is provided with SWVR. To use this file, click **Start**, then **Programs**, then **Compaq StorageWorks Virtual Replicator**, and then **Management Console**.

network disk name

A name that you choose for a network disk when you serve it. This is the name that the client computer or cluster uses to identify the disk when connecting to it.

Do not confuse this name with the volume label that you see in Windows NT Explorer. When the client connects to the disk, the client sees the volume label in Explorer, not the network disk name.

The network disk name is used only by SWVR software. If you want to change the network disk name, disconnect the client, stop serving the disk, and then serve it again, specifying a different network disk name.

node

A computer that is a member of a cluster.

online backup

The ability to back up your data without interrupting your users, and without compromising the consistency of your backup. While you are backing up a disk, your users can continue reading and writing it.

parent

When you create a snapshot, the disk that you take the snapshot of is called the parent.

pool

A collection of one or more physical disks or controller-based fault-tolerant disk sets that are used to store virtual disks and snapshots.

possible owners

Every cluster resource has a list of one or more possible owners. These are the nodes in the cluster that are capable of hosting the resource.

preferred node

The preferred owner of a cluster group is the node where you prefer the group to reside. When the preferred node fails, the group fails over to another node in the cluster, but when the preferred node comes back online, the group automatically fails back to it.

A group does not fail back if it does not have a preferred owner. You may choose not to select a preferred owner because it may not matter where the group resides, for example if all the nodes in the cluster are equally capable of handling the load required to run the group.

quorum

Voting mechanism used in a cluster to guarantee that the information necessary for recovery is maintained consistently between all cluster members. The mechanism involves a special disk called a quorum disk.

quorum disk

A disk used to store the cluster configuration information necessary for recovery of the cluster. The quorum disk is owned by only one cluster node at any point in time. When the node that owns the quorum disk fails, the surviving nodes arbitrate to take over its ownership.

Redundant Array of Inexpensive Disks (RAID)

A storage device that uses two or more magnetic or optical disks working in tandem to increase performance and provide various levels of error recovery and fault tolerance. RAID can be implemented in software using standard disk controllers, or it can be designed into the disk controller itself.

router

An intelligent hub that routes network traffic to a designated channel.

SCSI

See Small Computer System Interface.

segment

The smallest unit that is copied during a copy-out operation. It is a fixed-sized unit of contiguous blocks. The default size of a segment is 32 KB.

Small Computer System Interface (SCSI)

A protocol for connecting devices from several classes of peripherals to a host system without requiring modifications to hardware and software.

snap-in

A management tool that runs in the Microsoft Management Console (MMC).

snapshot

A disk in a pool that was created by making a virtual copy of another disk, called the parent disk.

When it is first created, the snapshot is a virtual copy of the parent disk, made at an instant in time. It has the same capacity and label, and it contains exactly the same data. It is as if you had a camera and took a picture of every single byte of data stored on the original disk at a single instant in time.

To applications and users, a snapshot looks like a normal disk. You can read and write it like a normal disk.

When you first create a snapshot, the snapshot takes up no disk space. It only starts to take up disk space when you modify the data stored on either the snapshot or the parent disk.

Snapshot Manager

A tool that you use to create and manage pools, virtual disks, and snapshots. It is an MMC snap-in. To run it, double-click the appropriate MMC console file.

A default console file is provided with SWVR. To use this file, click **Start** bar, then **Programs**, then **Compaq StorageWorks Virtual Replicator**, and then **Management Console**.

snapshot name

A name that you supply when you create a snapshot.

Do not confuse this name with the volume label that you see in Windows NT Explorer. The snapshot inherits the same volume label as its parent.

The snapshot name is used only by SWVR software. You cannot change it once you have created the snapshot.

standalone computer

A computer that is not in a cluster.

storage pool

See pool.

storage unit

A physical disk or controller-based fault-tolerant disk set whose space is donated to a pool. We use the term storage unit, instead of disk, to help distinguish it from the virtual disks that you create in the pool on which to store your data.

virtual disk

A disk in a pool that was created as an empty disk of a specified capacity. To your applications and users, it looks just like a normal disk. You map a drive letter to it, partition and format it like a normal disk, and then read and write it like a normal disk. But a virtual disk has one special feature – you can make virtual copies of it in a matter of seconds. The copies are called snapshots.

virtual disk name

A name that you choose when you create a virtual disk.

Do not confuse this name with the volume label that you see in Windows NT Explorer. By default, they are the same – when you use **Snapshot Manager** to format a virtual disk, the volume label is automatically set to the virtual disk name. But if you subsequently change the volume label, for example using Windows NT Explorer, the virtual disk name does not change.

The virtual disk name is used only by *Virtual Replicator* software. You cannot change it once you have created the virtual disk.

Wolfpack

The internal name used by the MSCS cluster project when the software was being developed.

Index

A

abbreviating commands 5-9

adding

- computers to registered client lists 7-3, 7-6, B-9
- IP subnets to client configurations 7-13
- snap-ins to MMC console files 5-3, 5-4
- snapshots to pools 6-6, A-21
- storage units to pools 6-12, A-12
- virtual disks to pools 6-5, A-34

Alpha serving network disks to Intel not supported 3-9, 9-9

application event log 9-2, 9-11

applications fail or hang 9-1, 9-5

available drive letters A-2, B-2

available storage units A-31

B

backup

- online 1-7, Glossary-4
- using snapshots for 6-16

Batch Scheduler Wizards 8-1

benefits of SWVR 1-6

C

capacity

- of pools 3-3, A-16
- of snapshots A-22, A-30
- of storage units A-31
- of virtual disks 6-5, A-43

case insensitivity

- in commands 5-9
- in names 3-9

Cluster Administrator

- when not to use 6-2, 6-5, 6-15, 7-9

cluster groups

- creating A-10, B-17, B-25
- defined Glossary-1
- deleting A-15, B-18, B-28

cluster names

- defined Glossary-1

cluster resources

- defined Glossary-1
- Physical Disk resource 7-3, B-16
- SCE Connected Disk resource 1-12, 9-11, B-25, B-28
- SCE Pool resource 1-12, 9-7, A-10, A-15
- SCE Served Disk resource 1-12, B-17
- SCE Storage Unit resource 1-12, A-15

clusters

- defined Glossary-1
- installing SWVR software on 5-1
- managing from standalone computers 5-1
- with nodes 100 miles apart 7-14

commands

- abbreviating 5-9
- how to use 5-6
- managing remote computers and clusters 5-8
- summary of 5-7
- upper or lower case in 5-9
- privileges required 5-9
- DRIVES A-2, B-2
- MANAGE A-4, B-4
- POOL A-8
- REGISTER B-8
- SERVE B-14
- SNAPSHOT A-21
- UNITS A-31
- USE B-23
- VIRTUALDISK A-33

computer names

- defined Glossary-2

computers, standalone

- defined Glossary-6

Index

- using to manage clusters 5-1
- connecting to network disks 7-8, B-24
- console files
 - adding snap-ins to 5-3, 5-4
 - default SWVR files 5-3
 - explained 5-2
- copy-out operations
 - caused by defragmenting disks 6-18
 - defined Glossary-2
 - explained 1-9
 - getting statistics on 6-23, A-18
 - predicting 3-4
- creating
 - cluster groups A-10, B-17, B-25
 - cluster resources 7-1, A-10, B-17, B-25
 - connections to served network disks 7-8, B-24
 - pools 6-2, A-9
 - served network disks 7-1, B-15
 - snapshots 6-6, A-21
 - virtual disks 6-5, A-34

D

- data mining and snapshots 1-7
- defragmenting virtual disks 6-18
- deleting
 - cluster groups A-15, B-18, B-28
 - cluster resources 7-1, A-15, B-28
 - connections to served network disks 7-21, B-28
 - pools 6-29, A-15
 - served network disks 7-22, B-18
 - snapshots 6-27, A-27
 - virtual disks 6-28, A-40
- Delspace
 - defined Glossary-2
 - explained 6-11
 - showing 6-11, A-30, A-43
- disconnecting from network disks 7-21, B-28
- disk space
 - available in a pool 6-9, 6-11, 9-4, 9-5
 - used by a snapshot 1-8, 3-4, 6-12, 6-18, 6-27, A-27, A-30
 - used by a virtual disk 3-4, A-43
- disks
 - see network disks

- see snapshots
- see storage units
- see virtual disks
- drive letters
 - mapping to snapshots 6-1, 6-7, 6-18, A-24
 - mapping to virtual disks 6-1, 6-6, 6-18, A-36
 - showing available A-2, B-2
 - unmapping from snapshots 6-20, A-26
 - unmapping from virtual disks 6-20, A-39

E

- error messages
 - failed to bring SCE Connected Disk resource online 9-2, 9-11
 - failed to bring SCE Pool resource online 9-7
 - failed write 9-5
 - lost connection to network disk B-26
 - lost delayed write 9-4
 - pool 70% full 6-10
 - pool 95% full 6-10
- event log 6-10, 9-2, 9-7, 9-11, B-26
- extended LAN
 - defined Glossary-2
- extension snap-ins for MMC 5-2

F

- failback
 - defined Glossary-2
- failover
 - defined Glossary-2
 - explained 1-12
- family
 - defined Glossary-3
 - explained 1-8, A-34
 - showing 6-25, 6-26, A-19, A-29, A-42
- FAT file system
 - not supported 3-3, 3-9
- formatting virtual disks 6-6, 6-19, 9-11, A-38
- free space in pools, monitoring 6-9

G

glossary of terms Glossary-1

I

I/O statistics

for network disks 7-19, B-21

for pools 6-23, A-18

I/Os to a network disk stall 9-1, B-26

installing SWVR on a cluster 5-1

Intel serving network disks to Alpha

not supported 3-9, 9-9

IP subnets

configuring network disk

clients 7-13

L

list of registered clients

see registered clients

local disks

defined Glossary-3

log. error 6-10, 9-2, 9-7, 9-11, B-26

M

Management Information Base 1-13

mapping drive letters

to snapshots 6-1, 6-7, 6-18, A-24

to virtual disks 6-1, 6-6, 6-18, A-36

MIB 1-13

Microsoft Management Console

see MMC

mirror sets, host-based

not supported 3-2, 3-3, 3-9

MMC

adding snap-ins 5-4

console files 5-2

defined Glossary-3

Results pane 5-3

Scope pane 5-2

tasks missing from 9-7

understanding 5-1

monitoring

I/Os to network disks 7-19, B-21

I/Os to pools 6-23, A-18

pool free space 6-9

.msc files 5-3

MSCS cluster

defined Glossary-3

multiple partitions

not supported 3-3, 3-9

N

names

case insensitivity of 3-9

cluster names Glossary-1

computer names Glossary-2

network disk names Glossary-4

snapshot names Glossary-6

virtual disk names Glossary-7

NDMgr commands

abbreviating 5-9

how to use 5-6

managing remote computers and
clusters 5-8

privileges required 5-9

summary of B-1

upper or lower case 5-9

DRIVES B-2

MANAGE B-4

REGISTER B-8

SERVE B-14

USE B-23

Network Disk Manager extension

snap-in 5-2

Network Disk Manager snap-in 5-2

network disks

benefits of 1-6

comparison with shares 1-11

connecting to 7-8, B-24

defined Glossary-3

disconnecting from 7-21, B-28

explained 1-4, 1-10

names of Glossary-4

performance of 1-12

restrictions on use of 3-9

servicing 7-1, B-15

showing information about 7-17,

7-19, B-20, B-30

stopping servicing 7-22, B-18

supported configurations 3-8

Index

- tuning tips 7-16
- using as cluster quorum
 - resource 7-14, B-27
- node
 - defined Glossary-4

O

- online backup 1-7, 6-16, Glossary-4

P

- parent
 - defined Glossary-4
 - explained 1-8
- partitioning virtual disks 6-19, A-38
- partitions, FAT
 - not supported 3-3
- partitions, multiple
 - not supported 3-3, 3-9
- Performance Monitor 1-14
- performance of network disks 1-12
- Physical Disk cluster resources 7-3, B-16
- planning SWVR configurations 3-1
- pools
 - adding storage units to 6-12, A-12
 - benefits of 1-6
 - capacity of 3-3, A-16
 - creating 6-2, A-9
 - defined Glossary-4
 - deleting 6-29, A-15
 - getting full 6-9, 9-4, 9-5
 - monitoring free space 6-9
 - reconstructing 9-12
 - showing information about 6-21, A-16
- possible owners
 - defined Glossary-4
- preferred node
 - defined Glossary-5
- privileges 5-9

Q

- quorum
 - defined Glossary-5
- quorum disk
 - defined Glossary-5
 - using a served network disk 7-14, B-27

R

- RAID
 - defined Glossary-5
- reconstructing a pool 9-12
- registered clients
 - adding 7-3, 7-6, 7-7, B-9
 - removing 7-3, 7-6, 7-7, B-11
 - showing 7-7, B-12
- removing
 - computers from registered client lists 7-3, 7-6, B-11
 - IP subnets from client configurations 7-13
 - snapshots from pools 6-27, A-27
 - virtual disks from pools 6-28, A-40
- restrictions
 - on disks you can serve 3-9
 - on network disk configurations 3-8
 - on use of virtual disks and snapshots 3-3
- Results pane of MMC 5-3

S

- SCE Admins group 5-9
- SCE Connected Disk resources 1-12, 9-11
 - creating B-25
 - deleting B-28
- SCE Pool resources 1-12, 9-7
 - creating A-10
 - deleting A-15
- SCE Served Disk resources 1-12
 - creating B-17
- SCE Storage Unit resources 1-12
 - deleting A-15
- Scope pane of MMC 5-2

- SCSI
 - defined Glossary-5
- security
 - controlling which clients connect to your served network disks 7-7, B-8
 - privileges for snap-ins and commands 5-9
 - SCE Admins group 5-9
- segments
 - defined Glossary-5
 - explained 1-10
- serving network disks
 - starting 7-1, B-15
 - stopping 7-22, B-18
- shares
 - comparison with network disks 1-11
 - using with network disks 1-10
- shortening commands 5-9
- showing
 - available drive letters A-2, B-2
 - available storage units A-31
 - information about network disks 7-17, 7-19, B-20, B-30
 - information about pools 6-21, A-16
 - information about snapshots 6-25, A-29
 - information about virtual disks 6-25, A-42
 - IP subnet configurations for clients 7-13
 - registered client lists 7-7, B-12
- size
 - see capacity
- snap-ins
 - adding 5-3, 5-4
 - defined Glossary-6
 - explained 5-2
 - extensions 5-2, 5-6
 - Network Disk Manager 5-2
 - privileges required 5-9
 - Snapshot Manager 5-2
 - using 5-1
- SnapMgr commands
 - abbreviating 5-9
 - how to use 5-6
 - managing remote computers and clusters 5-8
 - privileges required 5-9
 - summary of A-1
 - upper or lower case 5-9
- DRIVES A-2
- MANAGE A-4
- POOL A-8
- SNAPSHOT A-21
- UNITS A-31
- VIRTUALDISK A-33
- Snapshot Manager snap-in 5-2, 9-7
- Snapshot Planner 3-4
- snapshots
 - and disk structure corrupt errors 9-8
 - benefits of 1-6
 - creating 6-6, A-21
 - defined Glossary-6
 - deleting 6-27, A-27
 - disk space used by 1-8, 3-4, 6-12, 6-18, 6-27, A-27, A-30
 - explained 1-8
 - mapping drive letters to 6-1, 6-7, 6-18, A-24
 - names of Glossary-6
 - restrictions on use of 3-3
 - showing information about 6-25, A-29
 - unmapping drive letters from 6-20, A-26, A-27
 - using for backups 6-16
- SNMP agent 1-13
- SNMP Support 1-13
- stalled I/Os to a network disk 9-1, B-26
- standalone computers
 - defined Glossary-6
 - using to manage clusters 5-1
- statistics
 - I/Os to network disks 7-19, B-21
 - I/Os to pools 6-23, A-18
- stopping serving network disks 7-22, B-18
- storage pools
 - see pools
- storage units
 - adding to pools 6-12, A-12
 - creating a pool from 6-2, A-9
 - defined Glossary-7
 - showing which are available A-31
 - showing which are in a pool 6-22, A-18
- stripe sets, host-based
 - not supported 3-3, 3-9
- system event log
 - failed to bring pool online 9-7
 - lost connection to network disk B-26
 - pool 70% full 6-10

pool 95% full 6-10
reconnected to network disk B-26

W

Wolfpack clusters Glossary-7

T

taking snapshots 6-6, A-21
Task Automation Wizards 8-1
TCP/IP 1-12
terms Glossary-1
testing applications 1-7
troubleshooting 9-1
tuning tips for network disks 7-16

U

unmapping drive letters
from snapshots 6-20, A-26, A-27
from virtual disks 6-20, A-39, A-40

V

virtual disks
benefits of 1-6
creating 6-5, A-34
defined Glossary-7
defragmenting 6-18
deleting 6-28, A-40
disk space used by 3-4, A-43
mapping drive letters to 6-1, 6-6,
6-18, A-36
names of Glossary-7
partitioning and formatting 6-6,
6-19, 9-11, A-38
restrictions on use of 3-3
showing information about 6-25,
A-42
unmapping drive letters from 6-20,
A-39, A-40
volume sets
not supported 3-3, 3-9