

Compaq SANworks™

Virtual Replicator Version 2.0 System Administrator's Guide

Third Edition (June 2000)
Part Number: AA-RENC-TE
Compaq Computer Corporation

Notice

© 2000 Compaq Computer Corporation.

Compaq, the Compaq logo, and StorageWorks Registered in United States Patent and Trademark Office and other jurisdictions. SANworks and the SANworks logo are trademarks or service marks of Compaq Information Technologies Group, L.P.

Microsoft, Windows, and Windows NT are registered trademarks of Microsoft Corporation. Pentium is a trademark of Intel Corporation in the United States.

All other product names mentioned herein may be trademarks or registered trademarks of their respective companies.

Confidential computer software. Valid license from Compaq required for possession, use or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

Compaq shall not be liable for technical or editorial errors or omissions contained herein. The information in this document is subject to change without notice.

The information in this publication is subject to change without notice and is provided "AS IS" WITHOUT WARRANTY OF ANY KIND. THE ENTIRE RISK ARISING OUT OF THE USE OF THIS INFORMATION REMAINS WITH RECIPIENT, IN NO EVENT SHALL COMPAQ BE LIABLE FOR ANY DIRECT, CONSEQUENTIAL, INCIDENTAL, SPECIAL, PUNITIVE OR OTHER DAMAGES WHATSOEVER (INCLUDING WITHOUT LIMITATION, DAMAGES FOR LOSS OF BUSINESS PROFITS, BUSINESS INTERRUPTION OR LOSS OF BUSINESS INFORMATION), EVEN IF COMPAQ HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. THE FOREGOING SHALL APPLY REGARDLESS OF THE NEGLIGENCE OR OTHER FAULT OF EITHER PARTY AND REGARDLESS OF WHETHER SUCH LIABILITY SOUNDS IN CONTRACT, NEGLIGENCE, TORT, OR ANY OTHER THEORY OF LEGAL LIABILITY, AND NOTWITHSTANDING ANY FAILURE OF ESSENTIAL PURPOSE OF ANY LIMITED REMEDY.

The limited warranties for Compaq products are exclusively set forth in the documentation accompanying such products. Nothing herein should be construed as constituting a further or additional warranty.

Printed in the U.S.A.

Compaq SANworks Virtual Replicator System Administrator's Guide
Third Edition (June 2000)
Part Number AA-RENKC-TE

Contents

About This Guide

Intended Audience	xi
Related Documentation	xi
Document Conventions	xii

Chapter 1

Introduction

Features	1-2
Storage Pooling	1-2
Virtual Disks	1-2
Snapshots	1-3
Network Disks	1-4
Online Volume Growth	1-5
Cluster Support	1-5
SNMP Support	1-5
Performance Monitor Support	1-6
Interoperability with other Storage Management Tools	1-6
User Interface Options	1-7
Virtual Replicator at Work	1-7
Benefits	1-9

Chapter 2

Planning

Pools, Virtual Disks, and Snapshots	2-2
Storage Units in Pools	2-2
Rules for Using Virtual Disks and Snapshots	2-3
Pool Capacity	2-4
Disk Space Requirements for Virtual Disks and Snapshots	2-4
Using the Snapshot Planner	2-5
Disk Space Requirements for Pools	2-5
Advantages and Disadvantages of Large Pools	2-6
Existing Backup Tools	2-8
Checking if your Backup Tool Records Information on the Disk	2-8
Supported Network Disks	2-9
Disks Serving as Network Disks	2-10
Naming Conventions	2-10
Uniqueness of Names	2-11

Chapter 3

Installation

Prerequisites	3-1
Components	3-2
Installing Virtual Replicator	3-2
Installing Adobe Acrobat Reader	3-3
Getting Help	3-3

Chapter 4

Getting Started with Virtual Replicator

Using the Microsoft Management Console (MMC) Interface	4-2
Understanding MMC Snap-ins	4-2
MMC Panes	4-3
Performing Virtual Replicator Tasks in MMC	4-4
Adding a Snap-In	4-5
Managing Remote Computers	4-8
Using the Command Line Interface	4-8
Using Commands to Manage Remote Computers	4-10
Abbreviating Commands	4-11
Using Uppercase or Lowercase in Commands	4-11
Administering Security and Privileges	4-12
Managing Other Domains	4-12
Using Cluster Resources	4-13

Chapter 5

Guided Tour

Create a Pool	5-1
Create a Virtual Disk in the Pool	5-6
Create a Snapshot of the Virtual Disk	5-10
Serve the Snapshot to a Remote Computer	5-14
Serving the Snapshot	5-14
Adding an MMC Snap-in	5-16
Connecting the Remote Computer to the Served Snapshot	5-18

Chapter 6

Managing Pools

Creating Pools	6-2
Understanding Pool Free Space	6-5
Monitoring Pool Free Space using Snapshot Manager	6-6
Understanding Delspace	6-8
Adding Storage Units to a Pool	6-9
Displaying Information about Pools	6-12
Deleting Pools	6-15

Chapter 7

Managing Virtual Disks

Creating Virtual Disks	7-2
Restoring Virtual Disks	7-3
Mapping a Drive Letter to a Virtual Disk	7-4
Unmapping a Drive Letter from a Virtual Disk	7-5
Formatting a Virtual Disk	7-6
Defragmenting Virtual Disks	7-7
Displaying Information about Virtual Disks	7-7
Using Virtual Replicator	7-7
Using My Computer	7-9
Deleting Virtual Disks	7-10
Growing Virtual Disks	7-11

Chapter 8

Managing Snapshots

Introduction to Snapshots	8-2
Creating Snapshots	8-3
Creating Snapshots Automatically	8-5
Mapping a Drive Letter to a Snapshot	8-5
Unmapping a Drive Letter from a Snapshot	8-6
Using Snapshots to do Backups	8-7
Backing up a Local Disk Using the Command Line	8-7
Backing up a Network Disk Using the Command Line	8-8
Using the Wizard to Create Snapshots for Backup	8-10
Using Snapshots to Restore Virtual Disks	8-10
Other Backup and Restore Solutions	8-11
Displaying Information about Snapshots	8-11
Deleting Snapshots	8-13

Chapter 9

Managing Network Disks

Introduction to Network Disks	9-2
Serving a Disk	9-4
Serving Virtual Disks and Snapshots	9-4
Serving other Disks	9-5
Controlling Which Clients Connect to the Disks	9-8
Connecting to a Served Network Disk	9-9
Changing Settings	9-13
Changing the IP Subnet Configuration	9-13
Changing the Client System for a Served Network Disk	9-14
Using a Served Network Disk as the Cluster Quorum Disk	9-15
Steps When the Cluster Already Exists	9-15
Steps If the Cluster Does Not Already Exist	9-16
Tuning Tips	9-16
Showing Information about Network Disks	9-17
Showing Information about Served Network Disks	9-17
Showing Information about Connected Network Disks	9-20
Disconnecting from a Served Network Disk	9-21
Stopping the Serving of a Network Disk	9-22

Chapter 10

Growing Storage Volumes Online

Online Volume Growth for Basic and Virtual Disks	10-2
Understanding Basic, Virtual, and Dynamic Disks	10-2
Preparing for Online Volume Growth of Basic or Virtual Disks	10-4
Adding Storage Units to Grow a Pool	10-5
Growing Virtual Disks	10-5
Growing Basic Disks	10-8
Preparing Basic Disks for Growth	10-8
Creating Adjacent Space	10-9
Using Virtual Replicator to Grow Basic Disks	10-10
Example of Online Volume Growth for Basic Disks	10-12
Using Windows Explorer to Grow Basic and Virtual Disks	10-22

Chapter 11

Automating Tasks

Task Automation Wizards Overview	11-2
Setting the Scheduling Method	11-2
Using Windows Scheduled Tasks	11-3
Scheduling Tasks	11-3
Viewing and Modifying Task Automation Schedules	11-4
Using Compaq Batch Scheduler	11-5
Scheduling Tasks	11-5
Viewing and Modifying Task Automation Schedules	11-5

Chapter 12

Troubleshooting

Application using a Network Disk Fails or Hangs	12-2
Lost Delayed Write Errors	12-4
Applications Return Failed Write Errors	12-5
Accidentally Delete a Virtual Disk	12-6
Snap-in and Commands Do Not Show All the Pools in a Cluster	12-7
Disk Structure Corrupt Errors when Accessing a Snapshot	12-7
Disk Structure Corrupt Errors when you Stop Serving a Disk	12-8
Pool Free Space Fell by More Than the Amount of Data Written	12-9
Cannot Connect to a Network Disk in Another Subnet	12-10
SCE Connected Disk Resource Fails to Come Online in a Cluster	12-10
Reformatting a Virtual Disk Does Not Work	12-11
Performance Monitor Shows No Network Disks	12-11
Reconstructing a Pool	12-11
Tracking Errors	12-12
SnapDisk Driver Events	12-12
Network Disk Driver Events	12-13
Virtual Replicator Lifeguard Events	12-14

Appendix A

SnapMgr Commands

Appendix B

SmartSnap Command

Appendix C

NDMgr Commands

Appendix D

Virtual Replicator Lifeguard Service

Appendix E

Virtual Replicator Policy Management

Glossary

Figures

Figure 1–1. A Storage Pool with Virtual Disks and Snapshots	1-4
Figure 1–2. Example of Virtual Disks and Snapshots in a Data Center.	1-8
Figure 2–1. Options for Serving or Connecting to Network Disks	2-9
Figure 8–1. Example of a New Snapshot’s Disk Space Usage	8-2
Figure 8–2. A Copy-out Operation.	8-3
Figure 8–3. Disk Space Usage After a Copy-Out.	8-3
Figure 8–4. Using a Snapshot to Back Up a Network Disk	8-8
Figure 9–1. Serving a Disk Over a Network	9-2

About This Guide

Welcome to *Compaq SANworks™* Virtual Replicator.

Intended Audience

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

- Management of Microsoft Windows NT or Windows 2000 servers and workstations
- Configuration and management of Microsoft Cluster Server (MSCS) clusters

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

Related Documentation


In addition to this book, we supply these sources of information:

Where	What
In the box	<i>A Commands</i> quick reference card. <i>An Installation</i> card.
On the CD-ROM	In the Documentation folder of the Virtual Replicator 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• PDF versions of this book and the reference cards, so you can print extra copies of them
On the web	Visit http://www.compaq.com/sanworks/

Where	What
On the Virtual Replicator Help menu	Full information about using Virtual Replicator.

Document Conventions

The documentation conventions used in this guide are shown in the following table.

Documentation Conventions	
Keys, Menu Options, Command Names	Boldface type indicates user input. These elements also appear in initial capital letters.
<i>Command parameters</i>	Italic type indicates variable command parameters. It is also used for emphasis.
System display	Courier type indicates a keyboard command or screen display data.
FILENAMES, DIRECTORY NAMES, and DRIVE NAMES	These elements appear in uppercase.
Note:	Notes contain information that might be of special interest to the user.
Important:	Important notes contain clarifying information or specific instructions.
 Caution	Caution notes contain information that you need to know to prevent losing data or damaging the software or hardware.

Chapter 1

Introduction

Compaq SANworks Virtual Replicator provides advanced, centralized storage management capabilities in both clustered and non-clustered Microsoft 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 1-2)
 - Storage Pooling (Page 1-2)
 - Virtual Disks (Page 1-2)
 - Snapshots (Page 1-3)
 - Network Disks (Page 1-4)
 - Online Volume Growth (Page 1-5)
 - Cluster Support (Page 1-5)
 - SNMP Support (Page 1-5)
 - Performance Monitor Support (Page 1-6)
 - Interoperability with other Storage Management Tools (Page 1-6)
 - User Interface Options (Page 1-7)
- Virtual Replicator at Work (Page 1-7)
- Benefits (Page 1-9)

Features

Virtual Replicator provides manageable storage. System administrators can use a choice of user interfaces 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 TB (terabyte).
- Make instant copies, called **snapshots**, of the virtual disks.
- Use the virtual disks and snapshots on the **local** computer.
- **Serve** the virtual disks and snapshots over the **network** to remote computers.

Virtual Replicator offers support for cluster configurations, SNMP, and Microsoft Windows NT Performance Monitor.

Storage Pooling

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 *Compaq StorageWorks* RAID array do. The Virtual Replicator software controls how data is stored on a virtual disk. See Chapter 6, "Managing Pools," for more information.

Virtual Disks

The virtual disks that you create in the pool perform and behave in exactly the same way as physical disks. That is, you can format and map drive letters to them, and read from and write to them, just like physical disks. In addition, you can install applications to a virtual disk, including cluster-aware applications, such as Microsoft Exchange.

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. If you have a terabyte database, you can combine several disks or StorageWorks RAID arrays in a single pool, and create a terabyte virtual disk that spans that physical storage. The size of the virtual disks you create can range from 10 MB to 1 TB, depending on free pool space and other limits set at the time of pool creation.

See Chapter 7, “Managing Virtual Disks,” for more information.

Snapshots

Virtual Replicator lets you make instant replicas (called **snapshots**) of virtual disks in a matter of seconds. Snapshots enable the instant creation of multipurpose virtual replicas of production data without having to physically copy the data. Snapshots function in exactly the same way as ordinary physical disks with read and write capabilities.

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. *Compaq SANworks* Virtual Replicator provides wizards to automatically schedule unattended creation, deletion, and backup of snapshots. See Chapter 11, “Automating Tasks,” 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 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 almost no disk space from its pool. 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.

NOTE: When using advanced applications, such as database products, administrators need to be aware of the application’s requirements for data consistency. Before snapshots are taken, these applications might need to be paused to ensure data consistency.

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.




Figure 1-1. A Storage Pool with Virtual Disks and Snapshots

See Chapter 8, “Managing Snapshots,” for more information.

Network Disks

In addition to using virtual disks and snapshots directly on your Windows NT servers, 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 file services or network shares.

See Chapter 9, “Managing Network Disks,” for more information.

Online Volume Growth

With this release, Virtual Replicator features the ability to increase storage capacity without disrupting operations on Windows 2000 systems. Normally, when you grow a RAID set, the operating system does not recognize the size change until you reboot. The Virtual Replicator Online Volume Growth feature directs the operating system to update the size of a physical or virtual disk without requiring a reboot. Online volume growth also can instruct the operating system to grow the on-disk partition information for a volume. This allows a volume to grow into unused space made available by either increasing the size of a disk or by deleting one or more adjacent partitions.

With online volume growth, you can increase storage capacity, as required by your users and applications, easily and with zero downtime. See Chapter 10, “Growing Storage Volumes Online,” for more information.

NOTE: The online volume growth feature is available only on Windows 2000 systems.

Cluster Support

Virtual Replicator operates in conjunction with Microsoft Cluster Server (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.

The Virtual Replicator management tools automatically create cluster resources to simplify cluster management. See “Using Cluster Resources” on Page 4-13 for more information.

SNMP Support

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 local management information based on requests from applications that are monitoring 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\SANworks Virtual Replicator*.MIB

Locate the MIB file 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.

Performance Monitor Support

Virtual Replicator also supports Windows NT Performance Monitor. The Performance Monitor shows the behavior of virtual disks and snapshots using logical disk performance counters, such as Disk Read Bytes/sec and Disk Write Bytes/sec. The counters are similar to those for physical disks.

In addition, Virtual Replicator adds its own objects to the list of objects that can be monitored. These objects are named SWVR_Pool, SWVR_ServedDisk, and SWVR_Snapshot. Using special counters, you can monitor performance activity specific to Virtual Replicator pools, network disks, and snapshots, for example, the Delspace of a snapshot.

Interoperability with other Storage Management Tools

Virtual Replicator operates jointly with other storage applications from Compaq, such as SANworks SecurePath and StorageWorks Command Console, and other vendors. Refer to the Virtual Replicator Release Notes for specific information about using other storage applications.

Virtual Replicator can also be used with many database and messaging applications, such as SQL Server, Oracle, and Exchange. Go to the Compaq SANworks website, <http://www.compaq.com/sanworks>, to read Best Practices White Papers discussing these applications.

User Interface Options

Virtual Replicator offers users the following choices for interacting with the software:

- Microsoft Management Console (MMC) interface
- Command line interface

Each of these interfaces provides you with full capabilities to monitor and manage your storage resources. The MMC and command line interfaces work in a Windows NT environment. See Chapter 4, “Getting Started with Virtual Replicator,” for more information on the interface options.

Virtual Replicator at Work

The following example illustrates various features of Virtual Replicator working together.

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.

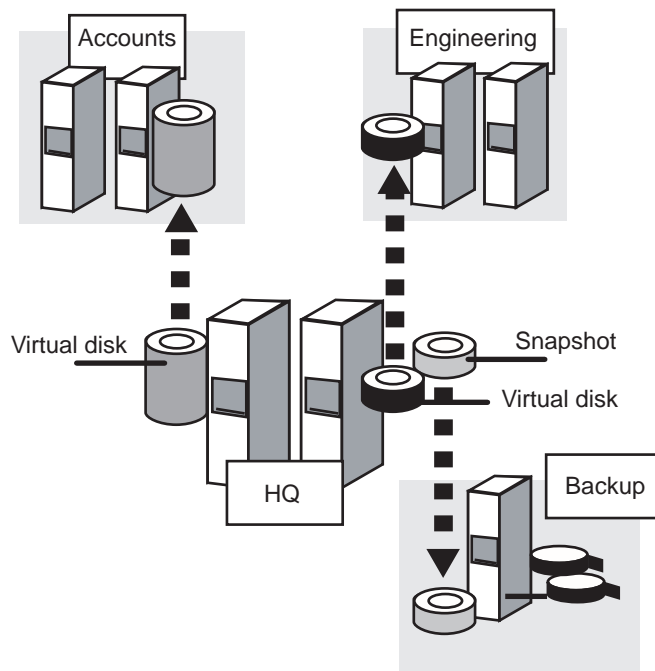


Figure 1–2. Example of Virtual Disks and Snapshots in a Data Center

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 administrator 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 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. Alternatively, on Windows 2000, the administrator can choose to grow the existing volume.

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.

Benefits

Here are some ways that typical enterprises can benefit from 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 each pool 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.

- **Add storage capacity and grow volumes without rebooting**

Virtual Replicator allows system administrators to understand when and where additional physical storage is needed. When a disk is added to a pool, Virtual Replicator directs the operating system to update the internal size of the pool. Then, without requiring a reboot, you can expand the sizes of virtual disks in the pool. Virtual Replicator also can be used to grow Windows 2000 basic disks.

- **Use virtual 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, dividing 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 in the following way:

- Create a 40-GB virtual disk for your Microsoft Exchange server.
- Create a 20-GB virtual disk for your web server.

The remaining 40 GB is spare space for your snapshots. 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, 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 as soon as you have created the snapshot.

- Use snapshots to keep copies of your data online and perform quick file restores

In addition to your standard backup procedures, snapshots can be used to perform quick restores of user data. 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.

- 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.

Snapshots eliminate the step of making a physical copy of your entire production data, and you do not have to take your data offline.

- Use snapshots with data mining applications

Data mining applications can populate your data warehouse by processing a snapshot of your live production data. This eliminates 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.

As an 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 Virtual Replicator tools to manage your environment remotely

The tools provided with Virtual Replicator let you manage local or remote computers or clusters. For example, you can create a virtual disk, map a drive letter to it, and format it, all from a remote computer.

■ 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. Create a snapshot at midday, then back up the snapshot to tape whenever you want. If there are no free tape drives, or if 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 snapshots to increase data availability

Virtual Replicator provides the option of making a snapshot copy of a virtual disk available via TCP/IP to another computer as a network disk.

Chapter 2

Planning

This chapter describes the information you need to consider when planning configuration for either a standalone computer or a cluster. The topics described in this chapter are:

- Pools, Virtual Disks, and Snapshots (Page 2-2)
- Storage Units in Pools (Page 2-2)
- Rules for Using Virtual Disks and Snapshots (Page 2-3)
- Pool Capacity (Page 2-4)
- Disk Space Requirements for Virtual Disks and Snapshots (Page 2-4)
- Disk Space Requirements for Pools (Page 2-5)
- Advantages and Disadvantages of Large Pools (Page 2-6)
- Existing Backup Tools (Page 2-8)
- Supported Network Disks (Page 2-9)

NOTE: To help plan the Virtual Replicator configurations, use the planning chart in the Documentation folder on the Virtual Replicator distribution CD-ROM. There is a blank chart to fill in and an example of a completed chart.

Pools, Virtual Disks, and Snapshots

Virtual Replicator lets you 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.

A family consists of a virtual disk, its snapshots, the snapshots of those snapshots, and so on. 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.

Storage Units in Pools

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

- Create pools from whole, unformatted disks, not from partitions or logical drives on a disk.
- 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 storagesets, or they should all be mirror sets, 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 sets. 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 2-3).

Another reason 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 must reconstruct the pool. You must 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, you should not use a standard single-spindle disk and a StorageWorks mirror set (RAID 1) in the same pool. Although the StorageWorks RAID storage set 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 weakest storage unit in this case is the standard disk.

Rules for Using Virtual Disks and Snapshots

You can use a virtual disk or snapshot in the same way as 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, 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, Virtual Replicator does not support software RAID.

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 storagesets, assume that the capacity of the pool is 90 GB for planning purposes. 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 Virtual Replicator.

Disk Space Requirements for Virtual Disks and Snapshots

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. To determine the size of a virtual disk, take into consideration the size of the pool and the number of snapshots you might want to take.

The disk space a snapshot uses can vary from 0 MB to the capacity of its parent virtual disk, depending on the rate at which your data changes, and 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). For example, if you have a virtual disk that is 10 GB in size, and *all* of the data on the virtual disk changes, a snapshot of the virtual disk could also grow to 10 GB.

If you have been making 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 an indicator of snapshot size. 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.

The next section, “Using the Snapshot Planner,” describes a Virtual Replicator tool that helps in planning disk space requirements.

Using the Snapshot Planner

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

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 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.

Snapshot Planner provides recommendations based on actual usage of your system. If the system environment changes, such as the addition of more applications or users, the values recommended by Snapshot Planner might need to increase or decrease.

To run Snapshot Planner, click **Start > Programs > Compaq SANworks Virtual Replicator > Snapshot Planner**.

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

NOTE: 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.

Disk Space Requirements for Pools

As with 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 might 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 might 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.



CAUTION: If a pool is full, you must not try to free up space by deleting files, because deleting files does not free up any space in the pool and can make the problem worse.

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 “Understanding Pool Free Space” on Page 6-5.

Advantages and Disadvantages of Large Pools

When deciding whether to create a large pool that contains several virtual disks or several small pools, each containing fewer virtual disks, keep these points in mind:

- 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 its virtual disks fail over to another node in the cluster. All applications must fail over with the pool at the same time.

To balance the load manually 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 storage sets. 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. Virtual Replicator does not do any load balancing across the different storage units in a pool.

- You can reduce the cost of disk space for snapshots by creating large pools that contain several virtual disks.

For example, assume that you want to create three 10-GB virtual disks. You want to use snapshots only to back up the disks. To ensure 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 the storage units available to you and decide how much disk space to give each application. If you have several small disks and an application that needs a large amount of disk space, you can create one pool out of the small disks.

Alternatively, if you have several applications that need a small amount of disk space and you have only a large disk, you should consider making a pool out of the one large disk and divide it into several small virtual disks.

If you do not want to use snapshots and are interested in saving disk space, you can combine disks into a pool. For example, you have 5 applications that each need 10 GB of disk space, but you have only 20-GB disks. You can 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 must take into account the disk space that the snapshots will use.

For example, you have a cluster that is running a Microsoft Exchange mail server and a Microsoft SQL Server 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 4 GB of disk space. You create a 20-GB pool for Exchange.
- 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 30-GB pool.

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 (such as using the archive bit) 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 subsequent incremental and differential backups do not work properly.

Checking if your Backup Tool Records Information on the Disk

Some backup tools record information about the backup by clearing the archive attribute 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 attribute, 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 to:

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 8-7.

Supported Network Disks

When planning your network disk configurations, remember that:

- In this release of 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.

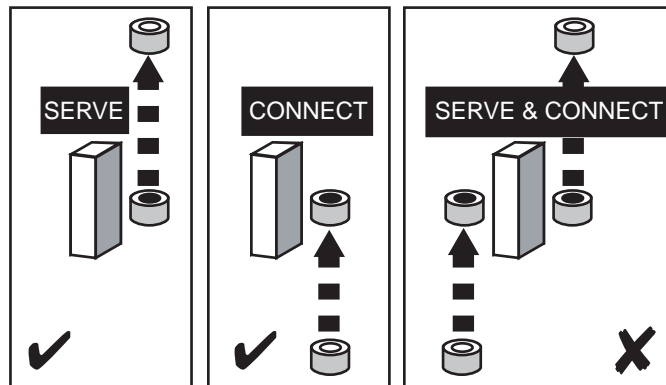


Figure 2-1. Options for Serving or Connecting to Network Disks

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 can connect to as many as 23 served network disks (drive letters D to Z) or serve any number of disks. A cluster can serve or connect to as many as 23 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 X86 or both be Alpha.

For example, an Alpha computer should not serve network disks to an Intel X86 computer.

Disks Serving as Network Disks

When planning which disks you want to serve as network disks, remember that:

- 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 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, 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, 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.

Naming Conventions

Every pool, virtual disk, and snapshot you create has a name. When you serve a disk, you provide a network disk name, which the remote client specifies when connecting to it.

NOTE: 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 the following 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 Virtual Replicator software shows which family each snapshot is in, but it does not show the parent relationship for snapshots of snapshots.

Uniqueness of Names

Every pool, virtual disk, and snapshot must have a different name. If you named a pool Pool1, you cannot create a snapshot named 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. If there is a cluster resource called Res1, do not create a pool named 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. If there is a cluster resource called Cluster1, do not name 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. If there is a cluster group named Group1, you cannot name the network disk Group1.

Chapter 3

Installation

This chapter includes the following topics:

- Prerequisites
- Components
- Installing Virtual Replicator
- Installing Adobe Acrobat Reader
- Getting Help

Prerequisites

Before you install Virtual Replicator V2.0, make sure the following software is installed:

- **Either** of the following operating systems:
 - Windows NT Version 4.0, with Service Pack 4, 5, or 6
 - Windows 2000, Professional, Server, or Advanced Server

NOTE: Of these versions of Windows 2000, only Advanced Server supports clustering.

- Microsoft Internet Explorer, Version 5.0 or greater

If you do not have Internet Explorer, you can download a free copy of the latest version from the Microsoft website at <http://www.microsoft.com/>.

The following hardware is required:

- A Pentium-based machine
- 128 MB of memory
- 40 MB of disk space for a full installation

Components

Version 2.0 of Virtual Replicator includes two installable components:

- Snapshot Planner

Before you install the Virtual Replicator component, you can use this tool to help predict how much disk space you need for your snapshots.

- Virtual Replicator

Use Virtual Replicator to create pools of storage, divide the pools into virtual disks, and create snapshots of the virtual disks. You can also serve any disks, including virtual disks and snapshots over the network.

Installing Virtual Replicator

1. Log on to a user account that has Administrator privileges.
2. Insert the Compaq SANworks Virtual Replicator installation CD into your CD-ROM drive. Installation begins automatically.

NOTE: If the installation does not start, you can double-click SETUP.EXE from the root directory on the CD.

3. To install the Snapshot Planner, click the component name on the first screen. Follow the instructions provided by the setup wizard until you are finished.
4. To install the Virtual Replicator, click the component name on the first screen.
5. In the License Agreement screen, click the button to accept the terms of the license agreement, then click **Next**.
6. In the Setup Type screen, choose the type of installation you want: Complete, Custom, or Management Only.

The default installation is Complete. Choose the Custom option if you want to select specific features. Choose the Management Only option if you want just the management tools that allow you to remotely administer Virtual Replicator.

7. Continue through the wizard screens until you have finished.

NOTE: Restart your computer after completing the installation if you receive an instruction to do so.

8. The next time you start your computer, make sure you log on to the **same user account** as in Step 1.

Installing in a cluster

The Virtual Replicator components are fully cluster aware. Be sure to install exactly the same Virtual Replicator options on all nodes of a cluster. For example, if you installed the Network Disk feature on one node, make sure you install it on all the nodes in the cluster.

Installing Adobe Acrobat Reader

PDF versions of this and several other documents are included in the Documentation folder on the installation CD. To view PDF files, you need to have Adobe Acrobat Reader installed. You can find an installation file for Acrobat Reader in the Documentation\Acrobat Reader folder. To install this software, use Windows Explorer to find the file, AR405ENG.EXE, on the installation CD and double-click it. Then follow the instructions provided by the installation wizard.

You can download a free copy of the latest version of Acrobat Reader from the Adobe website <http://www.adobe.com>

Getting Help

If you have a problem with Virtual Replicator, you can get further information and other help in the following locations.

Compaq Website

The Compaq website has information on this product. You can access the Compaq SANworks website at <http://www.compaq.com/sanworks>.

Compaq Authorized Reseller

For the name of your nearest Compaq Authorized Reseller:

- In the United States, call 1-800-345-1518.
- In Canada, call 1-800-263-5868.

Elsewhere, see the Compaq website for locations and telephone numbers.

Chapter 4

Getting Started with Virtual Replicator

This chapter describes how to use the Microsoft Management Console (MMC) and command-line interfaces to manage storage resources on standalone and cluster computers. The following topics are described in this chapter:

- Using the Microsoft Management Console (MMC) Interface (Page 4-2)
 - Understanding MMC Snap-ins (Page 4-2)
 - Adding a Snap-In (Page 4-5)
 - Managing Remote Computers (Page 4-8)
- Using the Command Line Interface (Page 4-8)
 - Using Commands to Manage Remote Computers (Page 4-10)
 - Abbreviating Commands (Page 4-11)
 - Using Uppercase or Lowercase in Commands (Page 4-11)
- Administering Security and Privileges (Page 4-12)
- Using Cluster Resources (Page 4-13)

Using the Microsoft Management Console (MMC) Interface

Microsoft Management Console (MMC) is the management environment from Microsoft. MMC lets you use different management tools and manage the local and remote computers from within the same window on your desktop. Microsoft and other software vendors are currently migrating their existing management tools into the MMC environment.

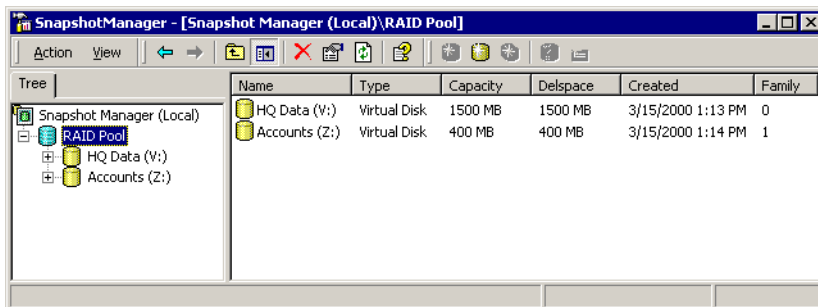
Understanding MMC Snap-ins

The management tools that you run in the MMC environment are called **snap-ins**. 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 serve virtual disks and snapshots.
Network Disk Manager	Manages network disks.

To start Virtual Replicator using MMC, click **Start > Programs > Compaq SANworks Virtual Replicator**, then either **Snapshot Manager** or **Network Disk Manager**.

MMC looks and works like Windows Explorer. The window is divided into two panes, as shown in the next 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 software vendor, and have, for example, 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 his or her own console file that contains all the snap-ins that they use, arranged in a tree view that suits the operator's particular needs.

Console files have the .msc extension. Virtual Replicator provides two default console files to get you started quickly, SNAPSHOTMANAGER.MSC and NETWORKDISKMANAGER.MSC. These files contain snap-ins to manage the local standalone computer or 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 4-5.

MMC Panes

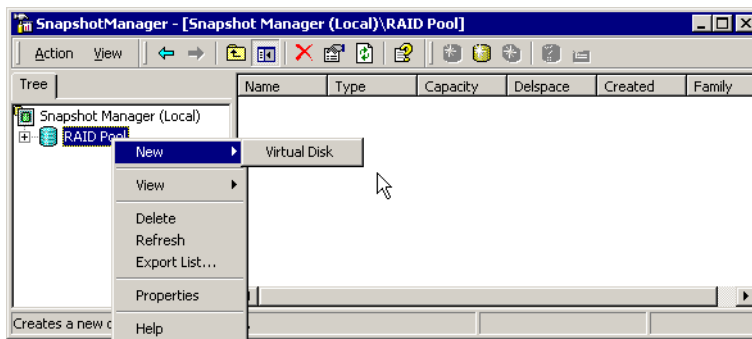
When you select an item in the left-hand Scope pane, the results of your selection are displayed in the right-hand **Results** pane. For example, when you select the Snapshot Manager snap-in in the left-hand Scope pane, the right-hand Results pane displays the pools on the standalone computer or cluster 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 or in an error state on the computer or cluster that the snap-in is managing.
A pool	All the virtual disks in the pool (for online pools).
A virtual disk	All the snapshots in the virtual disk's family (for online pools).
A snapshot	Nothing (the Results pane is empty).

When you select this in the left-hand Scope pane	The right-hand Results pane shows
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 are currently connected to the computer or cluster that the snap-in is managing.

Performing Virtual Replicator Tasks in MMC

The most common way to perform any Virtual Replicator task is to right-click an object in the left-hand Scope pane. This action opens a menu, as shown in the next screen capture.



Right-clicking a pool, virtual disk, or snapshot reveals different task menus.


In addition, all the options available on the right-click menus are also available on the MMC **Action** and **View** menus. For example, in the previous screen, instead of using the right-click menu, you could have clicked **Action**, and then **New > Pool**. These actions are also available as buttons on the toolbar.

Adding a Snap-In

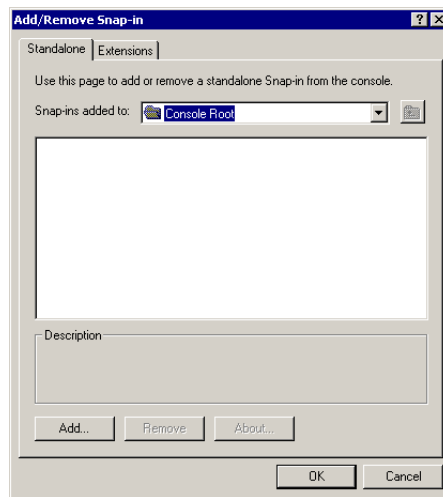
The two default snap-ins provided by Virtual Replicator, **Snapshot Manager** and **Network Disk Manager**, manage the local computer. You can create additional console files for managing other computers by adding one or more snap-ins to the MMC console. The new console file gives you a permanent mechanism for managing the remote machine(s). See the MMC Help for more information about creating a custom console file.

To add a snap-in to your MMC console:

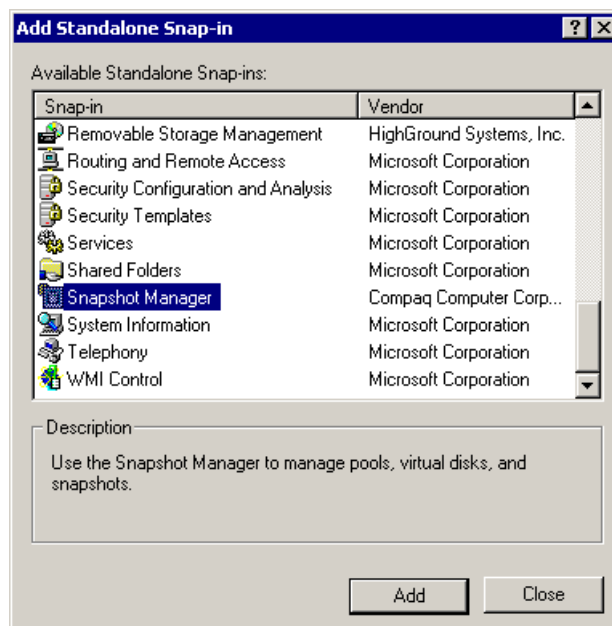
1. Click **Start**, then **Run**.
2. In the **Open** text box, type **mmc**.
3. On the MMC **Console** menu, click **Add/Remove Snap-in**.



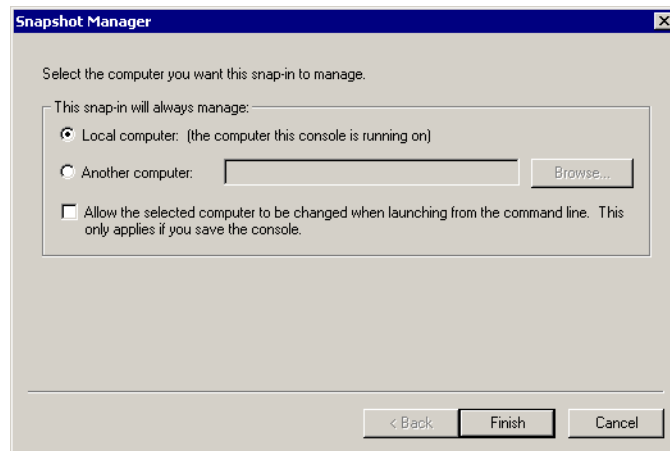
4. On the Add/Remove Snap-in screen, use the top list box to select the folder to which you want to add the snap-in. For example, select the **Console Root** folder.



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



6. Select the snap-in that you want to add. Select **Snapshot Manager** if you want to manage pools, virtual disks, and snapshots on the specified computer. Select **Network Disk Manager** if you want to manage network disks on the specified computer. Click **Add**. The screen for selecting a computer appears.



NOTE: 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.

7. 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**.

NOTE: 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.

8. Click **Close** to close the **Add Standalone Snap-in** dialog box.
9. To add a Snapshot Manager snap-in that can also **serve** the virtual disks and snapshots it manages, select the Network Disk Manager extension snap-in.

To do this, click the Extensions tab on the **Add/Remove Snap-in** dialog box, and check the box to the left of Network Disk Manager Extensions, then click OK.

10. Click **OK** to close the **Add/Remove Snap-in** dialog box.

11. Save the console file by clicking **File > Save As** and giving the console file a name.

The new console file runs in MMC, not in Snapshot Manager or Network Disk Manager.

Managing Remote Computers

You can use Virtual Replicator to manage both local and remote computers and clusters. By default, the Snapshot Manager and Network Disk Manager snap-ins manage the local standalone computer or cluster. Using the snap-ins, you can also manage a remote standalone or cluster computer if it has Virtual Replicator installed. The procedure below is functionally the same as using the MANAGE command (see Page A-4 and Page C-4 for more information).

To manage another standalone or cluster computer that has Virtual Replicator installed:

1. Click **Start > Programs > Compaq SANworks Virtual Replicator > Snapshot Manager** or **Network Disk Manager**.
2. In the left-hand Scope pane, select the **Snapshot Manager** or **Network Disk Manager** snap-in.
3. Right-click **Connect to another computer**.
4. In Windows NT 4.0, click **Another computer** to indicate the computer you want to manage. In the next dialog box, either type the name of the computer or cluster or browse for it on the network, and then click **Finish**.
5. In Windows 2000, in the **Select Computer** dialog box, either type the name of the computer or cluster or browse for it on the network, and then click **OK**.

You are connected and managing the remote computer. The console window shows all of the pools, virtual disks, and snapshots available on the remote computer. The connection is active until you close the snap-in. To set up a permanent mechanism for managing one or more remote computers, use MMC to create a custom console file by adding snap-ins for each computer (see “Adding a Snap-In” on Page 4-5).

Using the Command Line Interface

Virtual Replicator provides a command line interface that lets you write batch jobs to automate tasks. For example, you can write batch jobs to back up your data. You can run the commands from any command line.

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

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

Using Commands to Manage Remote Computers

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 application's command prompt or the Windows NT prompt.

At the application's 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.

At the Windows NT command prompt

If you are 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
    Available drive letters: GHKMNPQRST

C:\> SNAPMGR DRIVES
    Available drive letters: DGHMNPQ
```

Abbreviating Commands

Abbreviate the commands to three letters or less but do not 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
```

Using Uppercase or Lowercase in Commands

Type the commands in uppercase, lowercase, or a combination of both. For example, these commands are all the same:

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

Administering 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 in 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 let users manage the Virtual Replicator software without giving them the broad privileges associated with the Domain Admins global group:

1. Create a global group called SWVR Admins. A domain administrator must create this group.
2. Add the users to the SWVR Admins global group.
3. Add the SWVR 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.

Managing Other Domains

You can also manage a computer in another domain if:

- 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.

Using Cluster Resources

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

NOTE: 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 default properties.

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.

Chapter 5

Guided Tour

This chapter provides a guided tour giving you an overview of the Virtual Replicator tools that manage pools and network disks. This tour tells you how to perform the following tasks:

- Create a Pool (Page 5-1)
- Create a Virtual Disk in the Pool (Page 5-6)
- Create a Snapshot of the Virtual Disk (Page 5-10)
- Serve the Snapshot to a Remote Computer (Page 5-14)

For the purposes of this tour, the tasks occur in the Microsoft Management Console (MMC) interface. It should take you about 15 minutes to read through the tour – you do not need to follow the steps on your own computer.

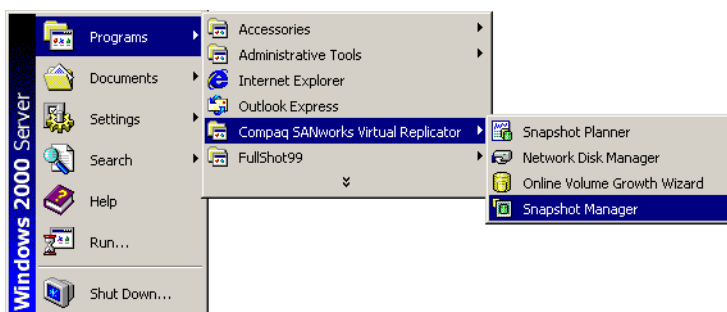
Create a Pool

This part of the tour creates a pool on the local HQ cluster using two 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 Chapter 6, "Managing Pools," for more information on creating pools.

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

1. Click **Start > Programs > Compaq SANworks Virtual Replicator > Snapshot Manager**.



The Microsoft Management Console (MMC) splash screen is briefly displayed. And then the Snapshot Manager console screen is displayed.

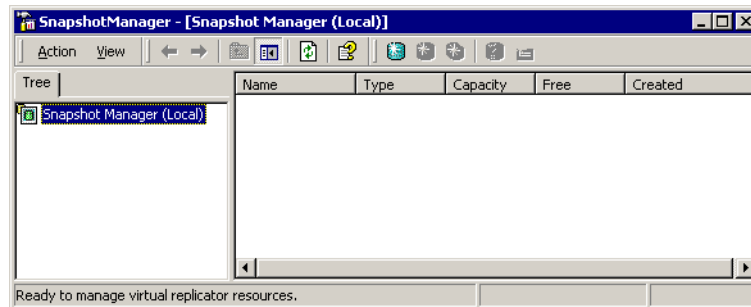


MMC is started and you have opened the **Snapshot Manager console file**, which helps you get started quickly.

MMC is similar to 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.

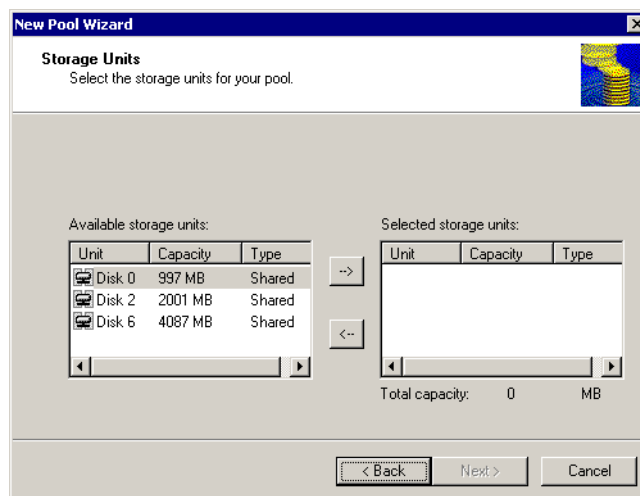
Click the items in the left-hand Scope pane to expand them. The right-hand **Results** pane displays information about whatever you have selected in the left-hand Scope pane.

2. Click **Snapshot Manager (Local)**, and then right-click **New** then **Pool**.



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

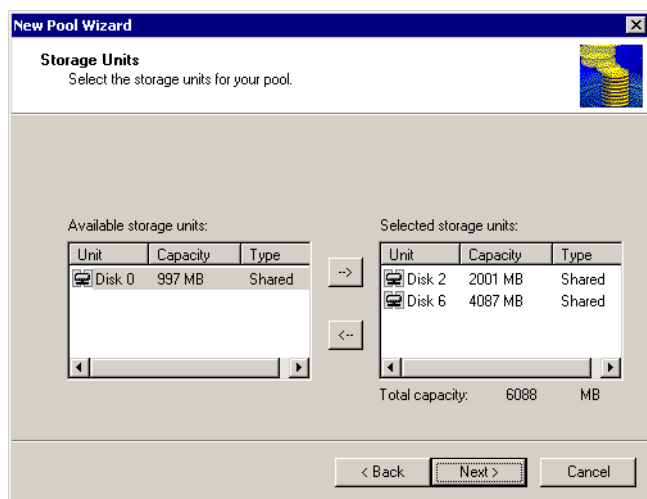
3. Click **Next** on the New Pool Wizard Welcome screen. The New Pool Wizard Storage Units screen is displayed.



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. In a cluster, select only shared disks.

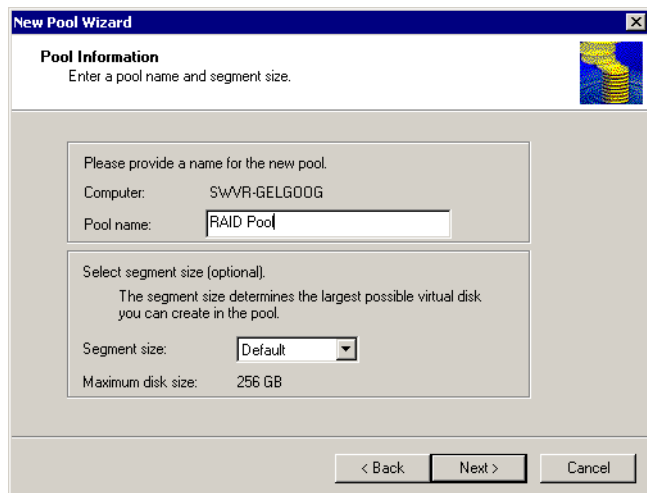
4. Select Disk 2 and Disk 6, then click the right arrow.

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.

5. Click **Next**. On the Pool Information screen, type the name that you want to give to the new pool. This example calls it RAID Pool.



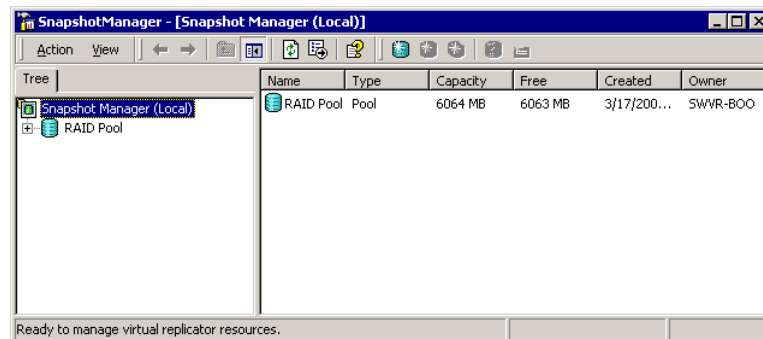
Then supply the **segment size** or accept the default. This value determines the largest possible virtual disk you can create in a pool. The default value is based on the size of the disks in the pool.

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



This screen tells whether you created the pool. As with all the Virtual Replicator final wizard screens, if the wizard fails, this screen provides an explanation.

7. Click **Finish** to close the final wizard screen. Now the management console looks like the following:

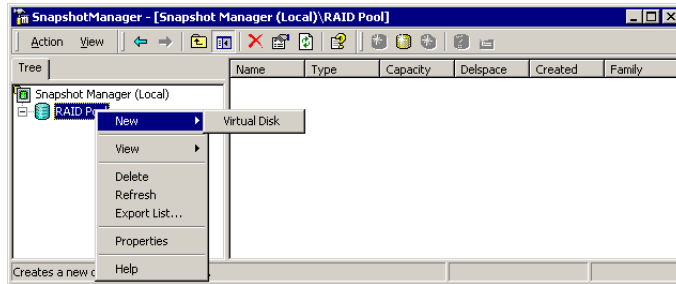


The new pool is listed. Notice that the pool's capacity is 1 MB more than its free space. That is because the capacity is rounded up and the free space is rounded down.

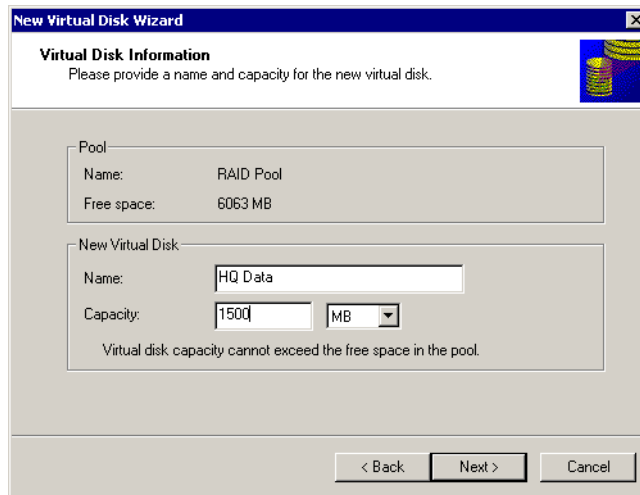
Create a Virtual Disk in the Pool

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.

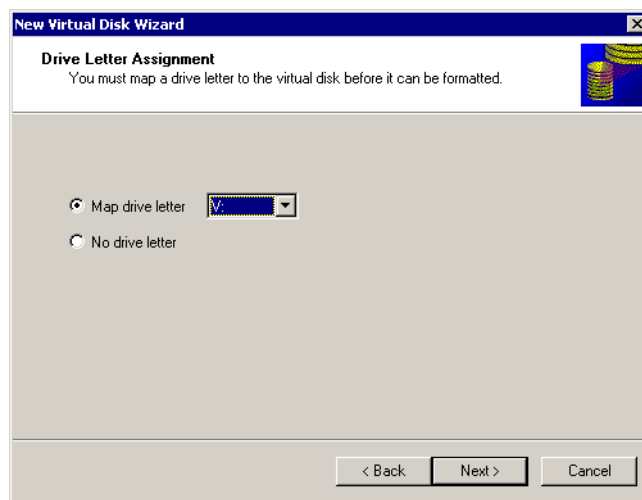
1. Select RAID Pool, and then right-click **New > Virtual Disk**.



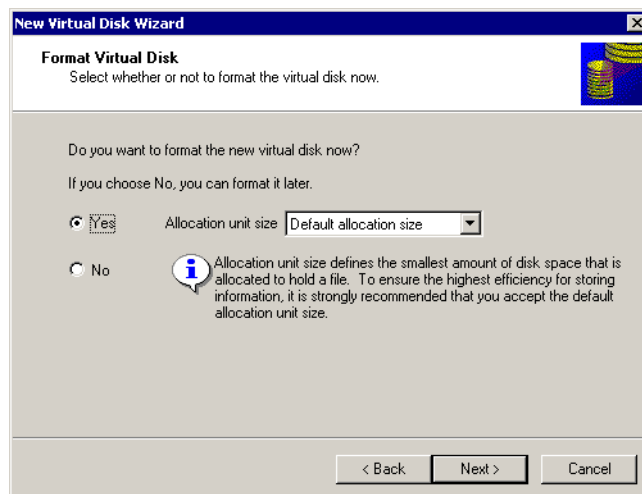
2. The welcome screen of the New Virtual Disk wizard is displayed. Click **Next** to start the wizard.
3. On the next screen, type the name of the new virtual disk and its capacity – this example calls it HQ Data and specifies 1500 MB – then click **Next**.



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**.

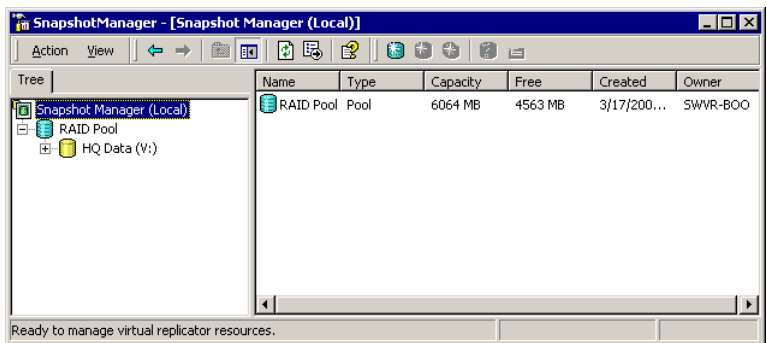


5. On the next screen, click **Yes** to format the virtual disk. Accept the default **Allocation unit size** to ensure that the disk is formatted for optimum storage efficiency. Click **Next**.



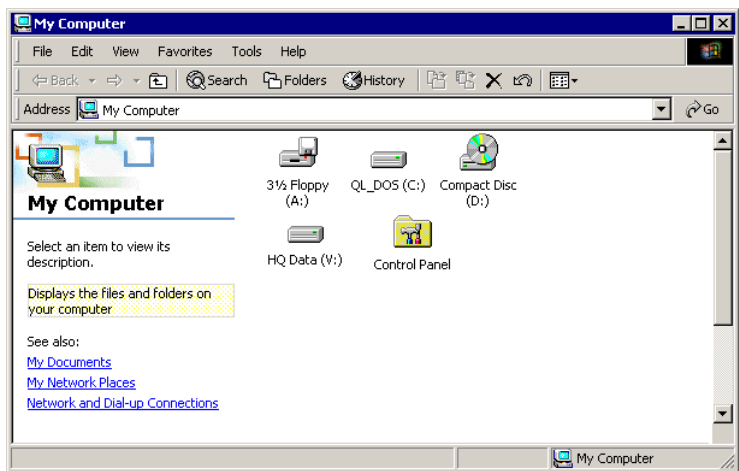
This creates a single partition on the disk, formats it with the NTFS file system, and gives it the same label as the virtual disk name (HQ Data).

6. The final wizard screen is displayed. Check that there are no error messages, then click **Finish** to close the wizard.
7. Select the **Snapshot Manager (Local)** folder.



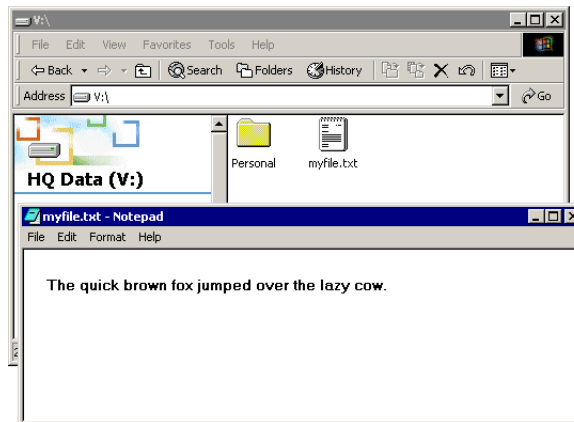
On the left, under RAID Pool, the new virtual disk, HQ Data, is indicated. The drive letter you mapped to it, V, is in brackets. On the right, notice that the pool's free space is now 4563 MB. It has dropped by 1500 MB, the capacity of the virtual disk you have just created.

8. Open My Computer.




The new disk appears with drive letter V, volume label HQ Data.

- Copy a folder called Personal and a file called Myfile.txt to the new disk:



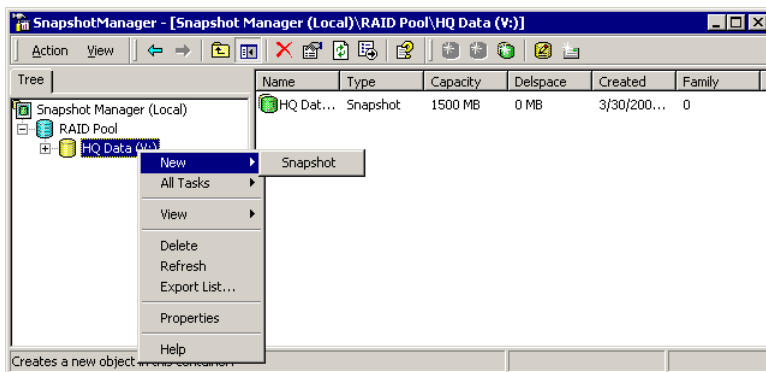
- Select the disk in My Computer, and look at the bottom status bar. The disk has a capacity of 1.46 GB (0.04 GB is being used by the file system for internal configuration data) and 1.45 GB of free space:



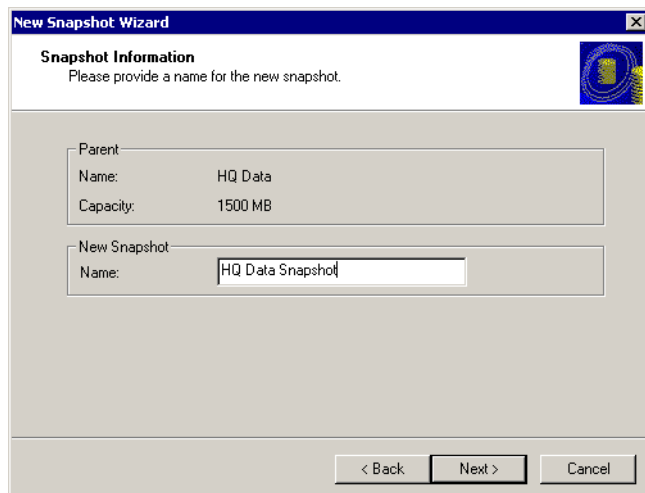
Create a Snapshot of the Virtual Disk

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 the **Snapshot Manager** console, select the virtual disk that you just created, HQ Data, then right-click **New > Snapshot**.

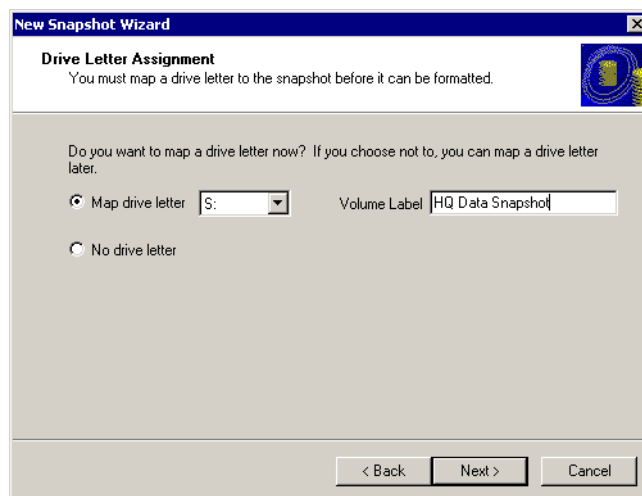


2. Click **Next** on the welcome screen to start the New Snapshot wizard.
3. Type a name for the new snapshot – HQ Data Snapshot – then click **Next**.




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.

The snapshot's Volume Label is, by default, the same as its parent disk. To avoid confusion between the two disks, change the snapshot's volume label to HQ Data Snapshot, then click **Next**.

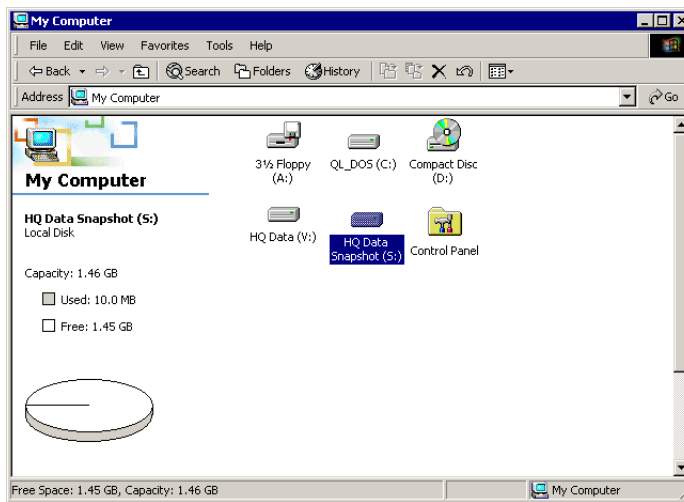


5. The final wizard screen is displayed. Verify that there are no error messages, then click **Finish** to dismiss the screen.

In the management console, the new snapshot is listed under its parent.

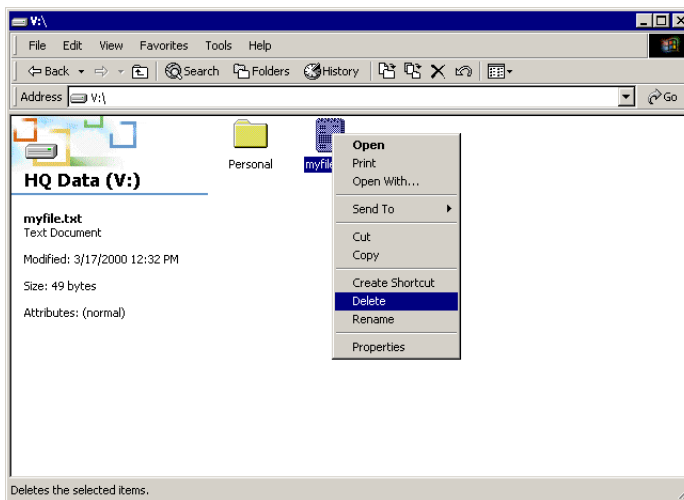


6. Open My Computer.

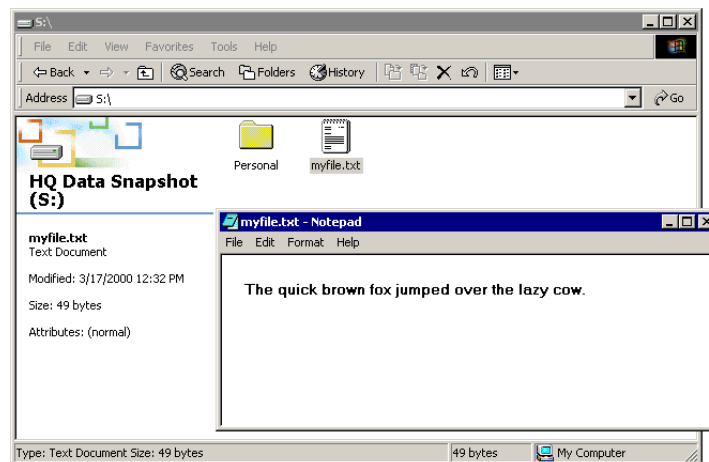


There is the new snapshot, HQ Data Snapshot. Select it and look at its free space and capacity on the bottom status bar. It has the same free space as its parent because it contains exactly the same files.

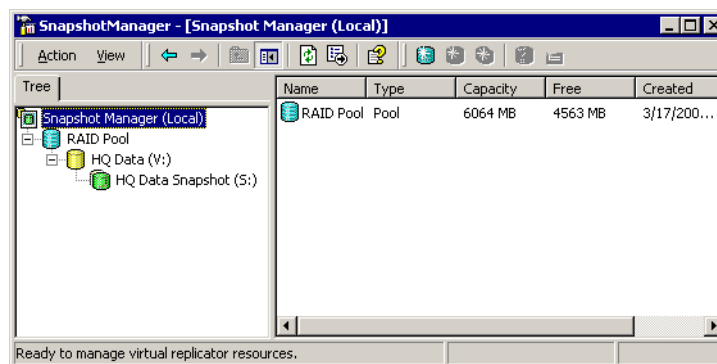
7. You can delete the file Myfile.txt on the parent disk, drive V.



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



8. Finally, go back to the Snapshot Manager console, and check the pool's free space. It is still 4563 MB; the snapshot is not using any disk space. The reason is that the only update that occurred on the disk was an update to the file system tables, which take very little space.



Serve the Snapshot to a Remote Computer

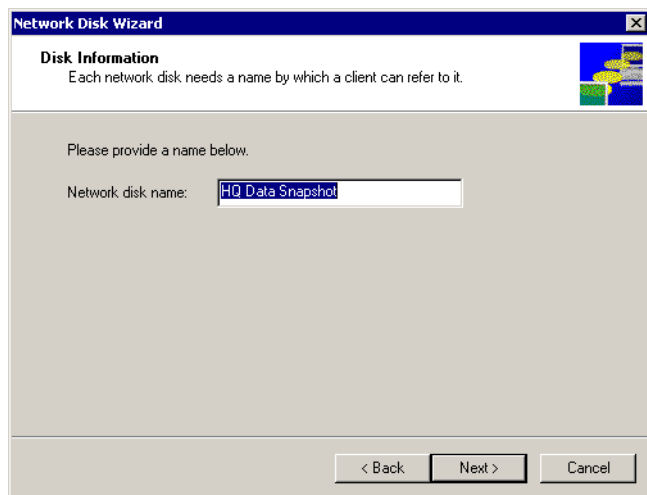
The final part of the guided tour serves the snapshot to the remote computer named SWVR-IKE. This involves three tasks:

- Serving the snapshot from the local computer.
- Adding a snap-in to the MMC console, so that you can manage the remote computer SWVR-IKE.
- Using this snap-in to connect SWVR-IKE to the served snapshot.

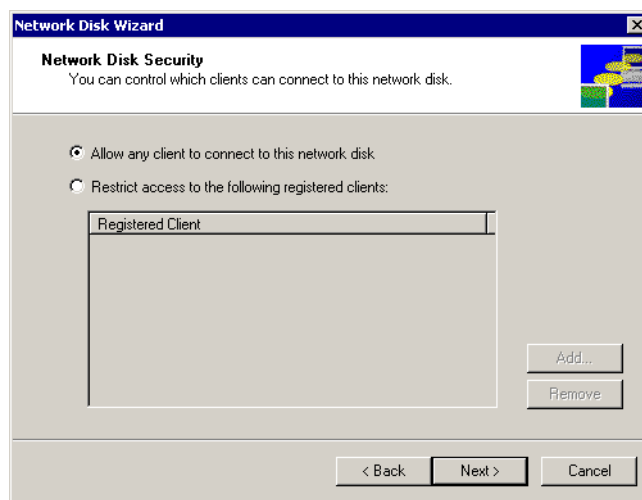
Serving the Snapshot

The first task is to serve the snapshot as a network disk.

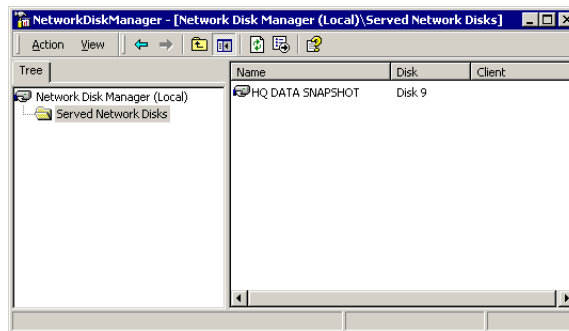
1. In the **Snapshot Manager** console, select the snapshot HQ Data Snapshot, and then right-click **All Tasks > Serve as Network Disk**.
2. The welcome screen of the New Network Disk wizard is displayed. Click **Next** to start the wizard.
3. On the Disk Information screen, click **Next** to accept the default network disk name (which is the same as the snapshot name, HQ Data Snapshot).



4. Click **Next** on the Network Disk Security screen to accept the default security settings, which allow any client computer to access the network disk.



5. The final wizard screen is displayed. Verify that there are no error messages, then click **Finish** to dismiss the screen.
6. Open the Network Disk Manager by clicking **Start > Programs > Compaq SANworks Virtual Replicator > Network Disk Manager**.
7. Select the Served Network Disks folder. The management console shows that HQ Data Snapshot is being served as a network disk.



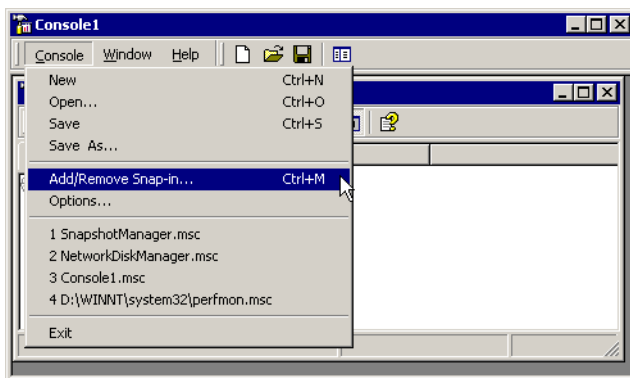
Notice that the name is in uppercase. Network Disk Manager displays all names in uppercase.

Minimize the Network Disk Manager console while you perform the next tasks.


Adding an MMC Snap-in

The second step is to add an MMC snap-in so that you can manage the remote computer SWVR-IKE.

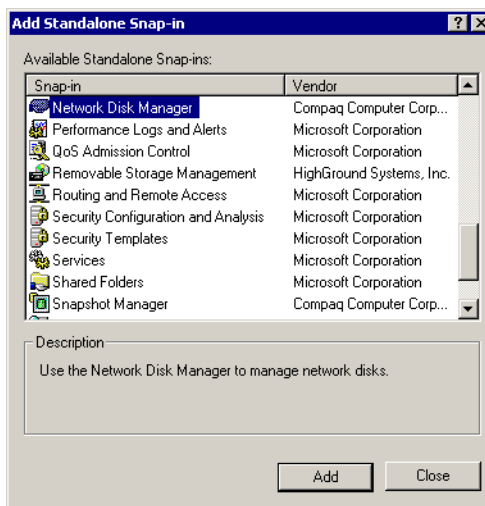
1. Click **Start**, then **Run**.
2. In the **Open** text box, type **mmc**.
3. On the MMC **Console** menu, click **Add/Remove Snap-in**.



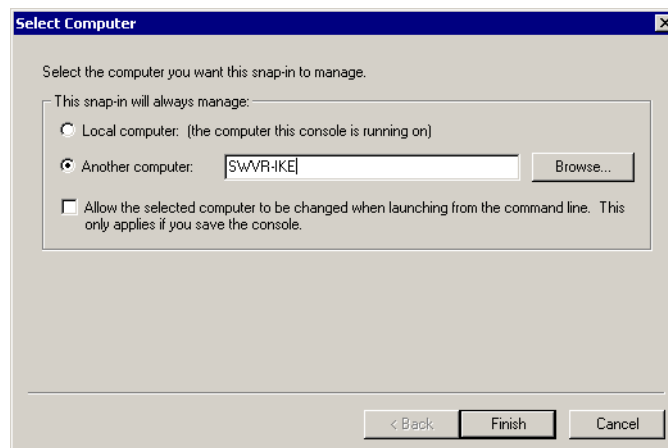
4. On the Add/Remove Snap-in screen, select the **Console Root** folder. Click **Add**.



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



6. On the Select Computer screen, click the **Another computer** button, then type in the name of the remote computer you want to manage, SWVR-IKE.

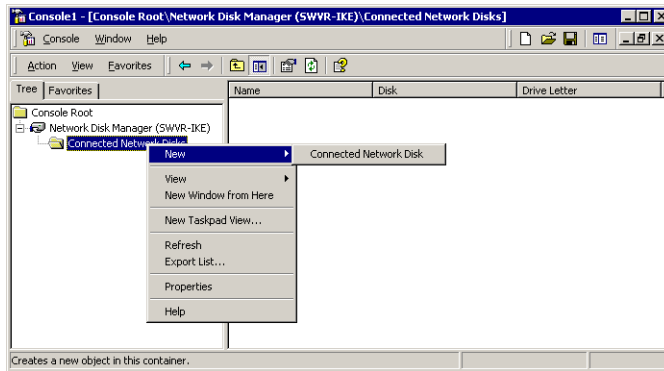


Click **Finish**, click **Close** on the Add Standalone Snap-in screen, and click **OK** on the Add/Remove Snap-in screen. You have added a Network Disk Manager snap-in to manage SWVR-IKE remotely.

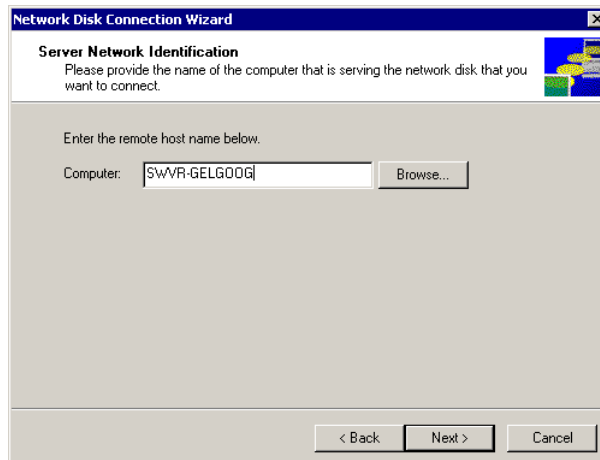
Connecting the Remote Computer to the Served Snapshot

The last step is to use the snap-in to connect SWVR-IKE to the served snapshot.

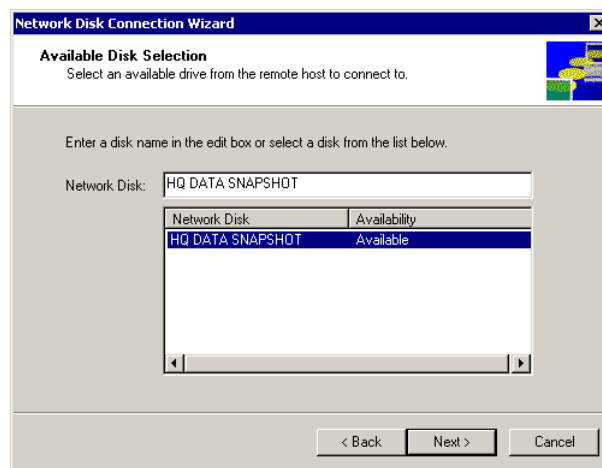
1. In the left pane of the SWVR-IKE console window (in MMC), click the plus sign in front of **Network Disk Manager (SWVR-IKE)** to see the **Connected Network Disks** folder. Select the folder, then right-click **New** then **Connected Network Disk**:



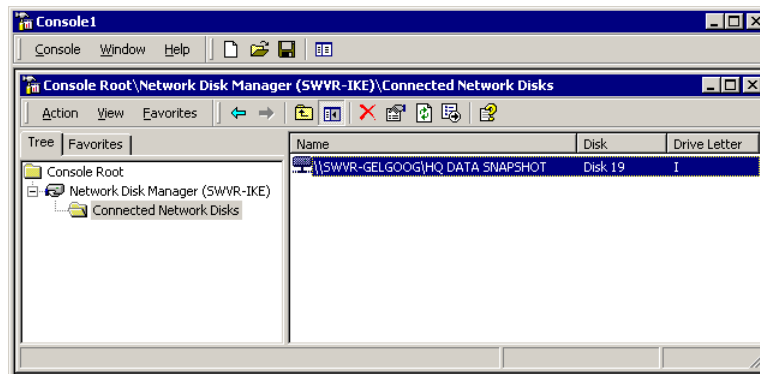
2. The welcome screen of the Network Disk Connection wizard is displayed. Click **Next** to start the wizard.
3. On the Server Network Identification screen, either type the name of the computer or cluster that is serving the network disk or select it from the browse list. In this example, the server's name is SWVR-GELGOOG. Type the name, and then click **Next**.



4. On the Available Disk Selection screen, select the network disk to connect to on the local computer. Select **HQ DATA SNAPSHOT**, then click **Next**.



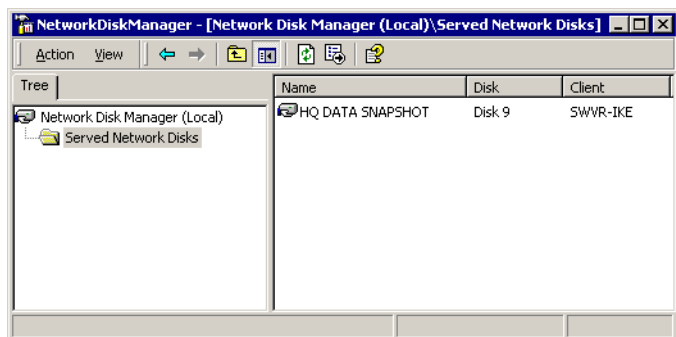
5. The final wizard screen is displayed. Verify that there are no error messages, then click **Finish** to dismiss the screen.
6. Select the **Connected Network Disks** folder under the snap-in **Network Disk Manager (SWVR-IKE)**.



The Name column on the right tells you that the network disk called HQ DATA SNAPSHOT is being served by SWVR-GELGOOG.

The Drive Letter column tells you that the drive letter **I** has been temporarily mapped to the disk on SWVR-IKE. This was the first available drive letter on SWVR-IKE. It is not persistent, so the next time SWVR-IKE restarts, a different drive letter may be mapped to the disk. You can make the drive letter persistent, on Windows NT 4.0 only, by using the Disk Administrator to set the drive letter.

7. Finally, look again at the Virtual Replicator Network Disk Manager console for the local machine (SWVR-GELGOOG). It shows that the client SWVR-IKE is connected to HQ DATA SNAPSHOT, which is being served by the local machine.



This concludes the Virtual Replicator guided tour. You should now be acquainted with the tasks of creating pools, virtual disks, and snapshots, and serving snapshots as network disks. These and other activities are explained in more detail in the chapters that follow.

Chapter 6

Managing Pools

Compaq SANworks Virtual Replicator gives you the ability to pool physical storage resources. After you create pools, you can then create virtual disks. Storage pooling and virtualization provide flexibility so that you can respond quickly to changing storage capacity requirements.

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

- Creating Pools (Page 6-2)
- Understanding Pool Free Space (Page 6-5)
- Monitoring Pool Free Space using Snapshot Manager (Page 6-6)
- Understanding Delspace (Page 6-8)
- Adding Storage Units to a Pool (Page 6-9)
- Displaying Information about Pools (Page 6-12)
- Deleting Pools (Page 6-15)

NOTE: 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 delete a pool in a cluster, not Cluster Administrator.

The tasks in the following sections require the Snapshot Manager snap-in, which is automatically added when you install Compaq SANworks Virtual Replicator. If the snap-in is not present, see “Adding a Snap-In” on Page 4-5.

Creating Pools

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.
- The largest capacity storage unit that can be included in a pool is 1 TB.
- 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 storagesets, or they should all be hardware mirror sets.
- In a cluster, use only the 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.

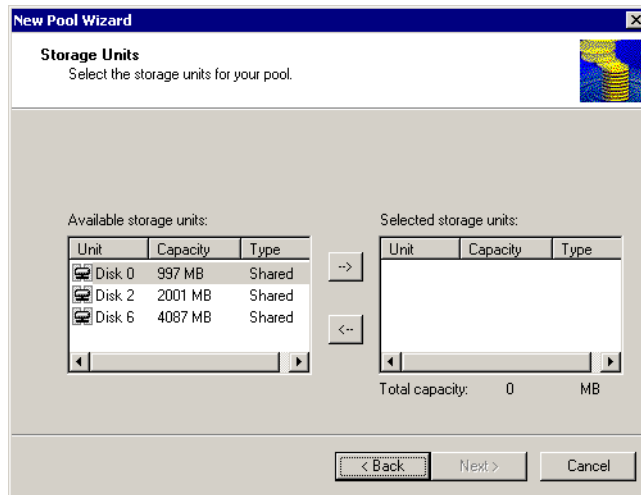
You can have any number of pools on a standalone computer or cluster. The number is limited only by the number of storage units that are available.

IMPORTANT: 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.

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. Click **Start > Programs > Compaq SANworks Virtual Replicator > Snapshot Manager**.
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.
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.

The wizard displays a list of the available storage units.



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.

CAUTION: The New Pool wizard may display disks that are part of a volume set, mirror set, or stripe set created using Disk Administrator's software fault tolerance feature. Do not select any of these disks or you could lose all of the data on the volume, mirror or stripe set. This is not a problem in clusters, because the MSCS software does not support these types of disks.

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.

NOTE: In a cluster, you cannot create a pool if any node in the cluster is down.

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	<p>The type of the storage unit:</p> <p>Shared: a disk in a cluster that is on a shared storage bus and can be seen by every node in the cluster.</p> <p>Local: either a physical disk that is attached to the computer, or a network disk.</p>
Total capacity	<p>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.</p> <p>The capacity of the pool will be about 10% less than the total capacity displayed.</p>

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

To select a storage unit, highlight it in the **Available storage units** list and then click the Right Arrow to move it to the **Selected storage units** list.

After you select the storage units you want to use, click **Next**.

NOTE: If any of the selected storage units is 1 TB or more, you need to remove it to continue with the wizard.

8. 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.

9. Choose a segment size or accept the default.

A **segment** is a fixed-sized unit of contiguous bytes of disk space. The specified size is used to define both the maximum virtual disk size you can create from the pool and the largest storage unit you can add to the pool. For example, choosing a segment size of 64 KB limits virtual disks and additional storage units in the pool to a maximum of 512 GB.

Segment size also determines the smallest unit that would be copied out to a snapshot when changes occur on a parent disk. See "Introduction to Snapshots" on Page 8-2 for more information on copy-outs.

Segment size options can range from 32 K to 128 K. The default value is the smallest of the available options. If you choose a different size, the maximum size of the virtual disk increases, but the amount of pool space that is consumed by snapshots might be greater.

Click **Next** to continue.

10. 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.

Creating a Pool 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.

Understanding 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 may fail and you may lose data. Use Snapshot Planner to help predict disk space requirements *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.

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, using the following methods:

- Use the Snapshot Manager snap-in to visually monitor pool free space, as described in “Monitoring Pool Free Space using Snapshot Manager” on Page 6-6.

- Check the System Event Log for Virtual Replicator warnings and errors about pool free space. When the free space falls to less than 30%, Virtual Replicator starts logging warnings every 5 minutes in the system event log to warn you that the pool is getting full. Then it logs errors every 5 minutes that say that the pool's free space is less than 5%. These records have **sddriver** in the source column. See "SnapDisk Driver Events" on Page 12-12 for a full description of these messages.
- If you want to gain further control over the monitoring of pool free space, use the Virtual Replicator Lifeguard service as described in Appendix D.

NOTE: You cannot use Windows Explorer to monitor pool free space because it does not recognize pools. Use only Virtual Replicator tools to monitor pool free space.

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. See "Adding Storage Units to a Pool" on Page 6-9, "Deleting Snapshots" on Page 8-13, and "Deleting Virtual Disks" on Page 7-10.

NOTE: 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.

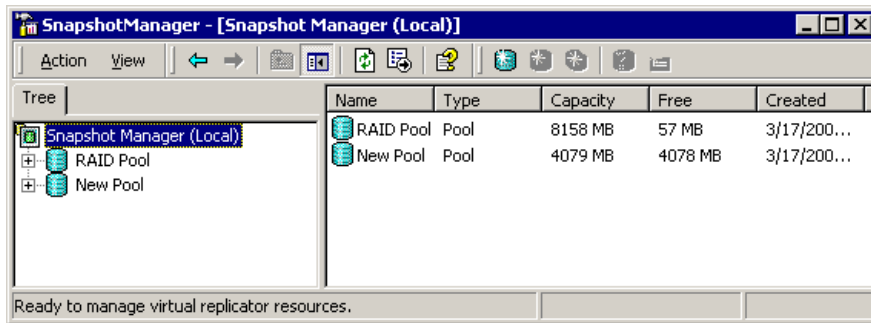
Monitoring Pool Free Space using Snapshot Manager

To monitor a pool's free space:

1. Click **Start > Programs > Compaq SANworks Virtual Replicator > Snapshot Manager**.
2. In the left-hand Scope pane, select the **Snapshot Manager** snap-in for the standalone computer or cluster that the pool is on.

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 the example below, the RAID Pool has a capacity of 8158 MB but it has only 57 MB of free space. It is almost full.



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

See “Adding Storage Units to a Pool” on Page 6-9, “Deleting Snapshots” on Page 8-13, and “Deleting Virtual Disks” on Page 7-10 for more information on these tasks.

It is important to remember that Windows Explorer does not recognize pools, and therefore, cannot be used to monitor free space in a pool. For example, the RAID Pool described above might have a virtual disk and two snapshots. If you were to use Windows Explorer to view the Properties pages for each of the disks, you would see information such as the following:

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

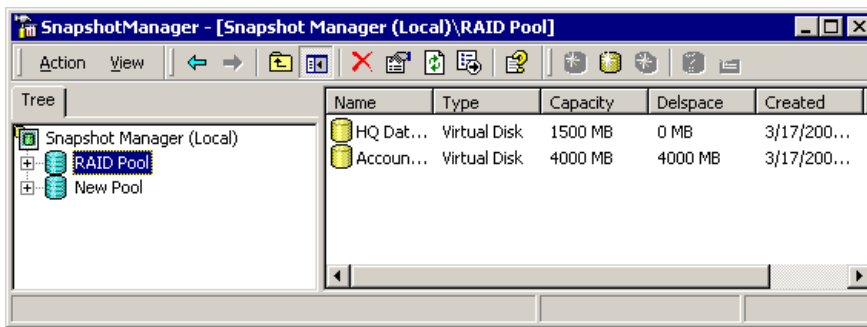
According to Windows Explorer, the combined free space on the three disks is 2,650 MB. However, the actual free space in the pool, as shown in the Snapshot Manager, is only 57 MB.

Likewise, deleting files has different effects in Windows Explorer and Snapshot Manager. For example, if you deleted files on drive Z: (VirtualDisk), you would see the free space rising in Windows Explorer. However, in Snapshot Manager free space would be unchanged because deleting files does not free up space in the pool. Accurate monitoring of pool free space is possible only with Snapshot Manager.

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.

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

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.

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 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.

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.
- Storage units can be standard single-spindle disks or controller-based fault-tolerant disk sets. They can have different capacities and be from different manufacturers.
- The largest capacity storage unit that can be included in a pool is 1 TB.

- The storage units 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 storagesets, you should only add StorageWorks RAID 5 storagesets.
- In a cluster, use only 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.

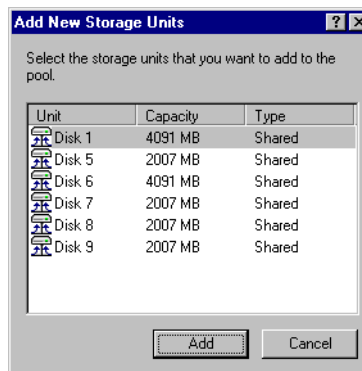
NOTE: In a cluster, you cannot add a storage unit to a pool if any of the nodes in the cluster are down.

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. Click **Start > Programs > Compaq SANworks Virtual Replicator > Snapshot Manager**.
5. In the left-hand Scope pane, select the **Snapshot Manager** snap-in for the standalone computer or cluster on which the pool resides.
6. Select the pool, then right-click **Properties** to display the Pool Properties dialog box.
7. 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.

8. Select the storage units you want to add, then click **Add**.

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

9. Click **OK** to dismiss the popup that tells you that the storage units have been added to the pool.

10. 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.

Adding Storage Units 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.

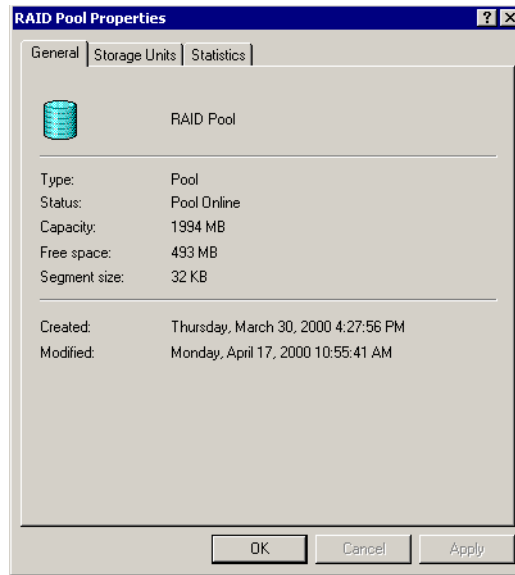
NOTE: Do not use Cluster Administrator to rename the storage unit resources, or remove the dependency of the pool resource on them, because you may not be able to access your data.

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.

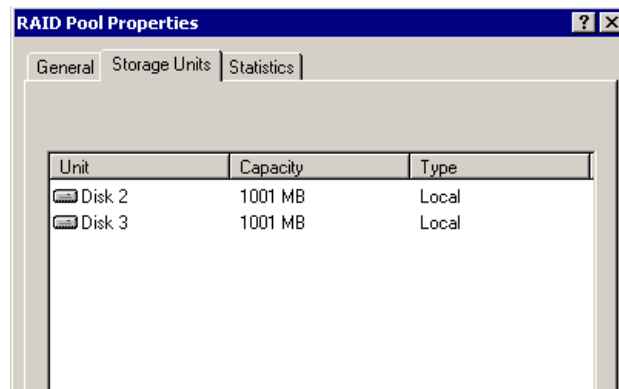
For more information, select the pool, then right-click **Properties**. The Properties page contains three tabs: **General**, **Storage Units**, and **Statistics**.



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

- The pool's status, for example, Pool Online.
- The pool's segment size. A segment is the smallest unit that is copied during a copy-out operation.
- The date it was last modified. This is when a storage unit was last added to it.

The **Storage Units** tab shows the storage units that are in the pool.



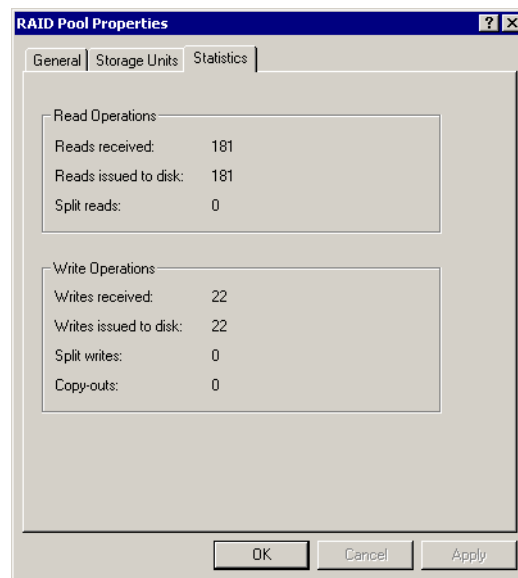
For each unit it shows the following information:

- 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.

The **Statistics** tab contains I/O statistics for 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
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.

Deleting Pools

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.

NOTE: You cannot delete a pool if it contains any virtual disks or snapshots.

To delete a pool:

1. Click **Start > Programs > Compaq SANworks Virtual Replicator > Snapshot Manager**.
2. In the left-hand Scope pane, select the **Snapshot Manager** snap-in for the standalone computer or cluster that the pool is on.
3. Select the pool, then right-click **Delete**.
4. Click **OK** to confirm that you want to delete the pool.

Deleting a 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.

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.

Chapter 7

Managing Virtual Disks

After you create pools (described in Chapter 6), you can then create virtual disks of a specific capacity. Virtual disks look and perform like physical disks, thereby providing flexibility for matching the disk space requirements of users and applications.

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

- Creating Virtual Disks (Page 7-2)
- Restoring Virtual Disks (Page 7-3)
- Mapping a Drive Letter to a Virtual Disk (Page 7-4)
- Unmapping a Drive Letter from a Virtual Disk (Page 7-5)
- Formatting a Virtual Disk (Page 7-6)
- Defragmenting Virtual Disks (Page 7-7)
- Displaying Information about Virtual Disks (Page 7-7)
- Deleting Virtual Disks (Page 7-10)
- Growing Virtual Disks (Page 7-11)

NOTE: 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, not Disk Administrator.

The tasks in the following sections require the Snapshot Manager snap-in, which is automatically added when you install Compaq SANworks Virtual Replicator. If the snap-in is not present, see “Adding a Snap-In” on Page 4-5.

Creating Virtual Disks

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 and is limited by the free space in the pool and the segment size of the pool. See “Creating Pools” on Page 6-2 for more information about segment size. You can create up to 8 virtual disks in each pool.

To create a virtual disk in a pool:

1. Click **Start > Programs > Compaq SANworks Virtual Replicator > Snapshot Manager**.
2. In the left-hand Scope pane, select the **Snapshot Manager** snap-in for the cluster or standalone computer that the pool is on.
3. Select the pool, and then right-click **New** then **Virtual Disk**.
4. The welcome screen of the New Virtual Disk wizard appears. Click **Next** to start the wizard.
5. Type a name for 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 or virtual disk 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 equal or exceed the free space in the pool, which is displayed by the wizard.

6. Map a drive letter to the new virtual disk (optional).

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 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” on Page 7-4.

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

Click **Yes** 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.

It is strongly recommended that you accept the default setting for **Allocation unit size**. This parameter defines the smallest amount of disk space that can be allocated to hold a file. Your operating system bases the default allocation unit on the size of the virtual disk. This minimizes the amount of lost space and the amount of fragmentation on the volume. Therefore, to ensure the highest efficiency for storing information, let the operating system choose the allocation unit size.

Click **Next** to proceed with creating a virtual disk.

If you do not want to format the disk now, click **No**, then click **Next**. To format the disk later, see “Formatting a Virtual Disk” on Page 7-6.

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

Creating a Virtual Disk 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. The virtual disks in a pool automatically appear on whichever node currently owns the pool resource.

Restoring Virtual Disks

When data on a virtual disk becomes corrupted, you can restore the virtual disk from your normal tape backup. However, a quicker method that minimizes downtime is to restore the virtual disk from a snapshot.

NOTE: To restore a failed virtual disk, you need to already have a snapshot of the virtual disk. See Chapter 8, “Managing Snapshots,” for information on creating snapshots.

Using snapshots, 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 a previously created snapshot of the virtual disk.
- **Method 2:** Assign a drive letter to a previously created 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 a previously created snapshot to the new virtual disk. Once the copy is completed, delete the old virtual disk and all of its snapshots.

Although Method 1 is the quickest way to make data available to users, it leaves the original virtual disk as an orphan and reduces performance because of additional 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.

You can schedule the Method 3 type of online restore process using the Restore From Snapshot wizard. See Chapter 11, "Automating Tasks," for a description of this wizard.

Mapping a Drive Letter to a Virtual Disk

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

To map a drive letter to a virtual disk:

1. Click **Start > Programs > Compaq SANworks Virtual Replicator > Snapshot Manager**.
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.
3. Select the virtual disk, and then right-click **All Tasks > Map Drive Letter**.

4. Select a drive letter from the pull-down list.

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. 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.

5. Click **OK**.

The drive letter is mapped to the virtual disk and then the disk is brought online.

Unmapping a Drive Letter from a Virtual Disk

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

To unmap a drive letter from a virtual disk:

1. Make sure that no one is accessing the virtual disk. You cannot unmap its drive letter if any files on it are open.
2. Click **Start > Programs > Compaq SANworks Virtual Replicator > Snapshot Manager**.
3. In the left-hand Scope pane, select the **Snapshot Manager** snap-in for the standalone computer or cluster that the virtual disk is on.
4. Select the virtual disk, and then right-click **All Tasks > Unmap Drive Letter**.
5. Click **Yes** to confirm the unmapping.

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

Formatting a Virtual Disk

When a virtual disk is created, it consists of a single partition that consumes the entire disk. If you did not format the virtual disk when you created it, you can use Snapshot Manager to format it later. Formatting can be used to wipe the virtual disk clean, or to change NTFS disk parameters, such as allocation unit size.

IMPORTANT: Formatting erases all the data on the virtual disk.

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” on Page 7-4.

You can also use Disk Administrator to format a virtual disk. When doing so, however, make sure that you select the NTFS file system. Virtual Replicator does not support the FAT file system.

IMPORTANT: Never use Disk Administrator to repartition a virtual disk.

To format a virtual disk:

1. Click **Start > Programs > Compaq SANworks Virtual Replicator > Snapshot Manager**.
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.
3. Select the virtual disk, and then right-click **All Tasks > Format**.
4. The Format dialog box shows the current settings for the virtual disk.

You cannot change the settings for **Capacity** or **File System**.

You can change the **Allocation unit size**, however, it is strongly recommended that you accept the default setting. This parameter defines the smallest amount of disk space that can be allocated to hold a file. Your operating system bases the default allocation unit on the size of the virtual disk. This default setting minimizes the amount of lost space and the amount of fragmentation on the volume. Therefore, to ensure the highest efficiency for storing information, let the operating system choose the allocation unit size.

You can also change the disk's current **Volume label**. The volume label is the name used by Windows to identify the virtual disk. If you change the label in the Volume label text box, the name appears next to the assigned drive letter in Windows Explorer. It also appears in the Windows Property page for the virtual disk. Changing the volume label does not change the name of the virtual disk in Virtual Replicator.

5. Click **Start** to begin formatting.
6. Click **OK** to confirm that you want to format the disk and delete all the data on it.
This formats the virtual disk with the specified settings.
7. Click **OK** to close the Format Complete dialog box.
8. Click **Close** to close the Format dialog box.

Defragmenting Virtual Disks

Before you defragment a virtual disk, 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. See “Introduction to Snapshots” on Page 8-2 for more information. To avoid unnecessary copy-out operations, delete all the snapshots on a virtual disk before defragmenting.

Displaying Information about Virtual Disks

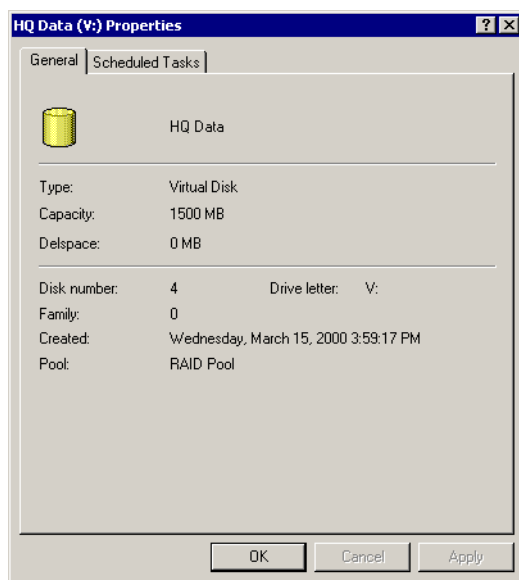
Information about virtual disks is available from both the Virtual Replicator Snapshot Manager and the Windows **My Computer** icon.

Using Virtual Replicator

To display information about a virtual disk using the Compaq SANworks Virtual Replicator software, select its pool in the left-hand Scope pane of the Snapshot Manager console window. 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	The family that the virtual disk is in.
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, then right-click **Properties**.



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

The **Scheduled Tasks** tab lists any tasks that have been scheduled for this virtual disk. See “Using Windows Scheduled Tasks” on Page 11-3 for more information.

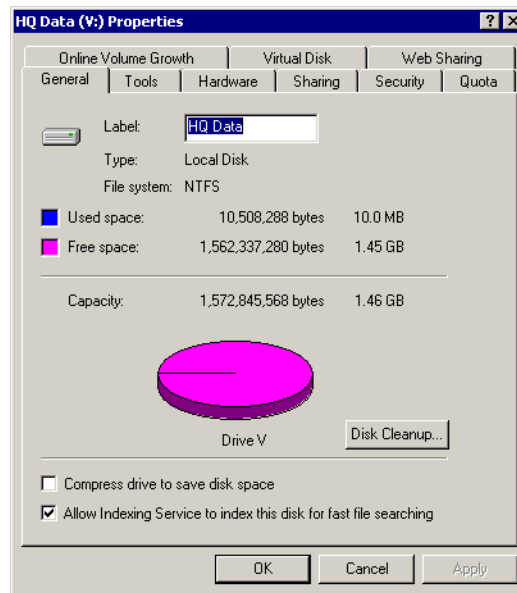
If the virtual disk is being served, there are **Network Disk Info** and **Network Disk Statistics** tabs, which show information such as the network disk name and the name of the client computer that is connected to it. For information on what these tabs display, see “Showing Information about Served Network Disks” on Page 9-17.

Using My Computer

To show information about a virtual disk using the My Computer icon:

1. Double-click **My Computer** on the Windows desktop.
2. Right-click the icon for the virtual disk, then click **Properties**.

The Windows **Properties** pages appear, as shown below.



The **General** tab shows the following information about the virtual disk.

Field	Description
Label	The Windows volume label of the virtual disk. You can edit the name in the text box. This will change the volume name in Windows, but not the name of the virtual disk in Virtual Replicator.
Type	The type of disk, for example, local or network.
File system	The file system for this volume, for example, NTFS (NT File System).
Used space	The amount of virtual disk space currently being used.
Free space	The amount of virtual disk space that is currently available.
Capacity	The total capacity of the virtual disk.
Drive	The drive letter to which the virtual disk is mapped.

The **Virtual Disk** tab shows the same properties that you see when you view the **Properties** page in Virtual Replicator (see Page 7-7).

Deleting Virtual Disks

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

NOTE: 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.

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. Click **Start > Programs > Compaq SANworks Virtual Replicator > Snapshot Manager**.

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.
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.

NOTE: When you delete a virtual disk, make sure to delete any scheduled tasks for the disk. Otherwise, if you reuse the old virtual disk name for another virtual disk, the old tasks will run on the new disk. Use the Windows Scheduled Tasks applet to delete scheduled tasks.

Growing Virtual Disks

Virtual Replicator provides an online volume growth feature that lets you increase the capacity of virtual disks without having to restart the computer. This feature is available only on Windows 2000. See Chapter 10, “Growing Storage Volumes Online,” for detailed information about Online Volume Growth.

Chapter 8

Managing Snapshots

After you create a virtual disk (described in Chapter 7), you can begin to take snapshots of it. A snapshot is an instant replica of a virtual disk. Snapshots are useful at such times as when you want to test new applications or when you need to restore data after a disk error or failure.

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

- Introduction to Snapshots (Page 8-2)
- Creating Snapshots (Page 8-3)
- Mapping a Drive Letter to a Snapshot (Page 8-5)
- Unmapping a Drive Letter from a Snapshot (Page 8-6)
- Using Snapshots to do Backups (Page 8-7)
- Displaying Information about Snapshots (Page 8-11)
- Deleting Snapshots (Page 8-13)

NOTE: 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 snapshots, not Disk Administrator.

The tasks in the following sections require the Snapshot Manager snap-in, which is automatically added when you install Virtual Replicator. If the snap-in is not present, see “Adding a Snap-In” on Page 4-5.

Introduction to Snapshots

A snapshot is a disk in a pool that is created by making a virtual copy of another disk, 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.

Initially, the snapshot does not use any disk space. It shares the same physical disk space as its parent. But if users create, modify, or delete the files or directories that are stored on either the snapshot or the parent disk, the data being modified must be physically copied beforehand. This operation is called a **copy-out**.

A snapshot starts to use disk space only when the parent or snapshot is modified and a copy-out is performed. The snapshot consumes disk space in units called segments. A segment is a fixed-sized unit of contiguous bytes of disk space. The segment size is set when you create the pool (see “Creating Pools” on Page 6-2). Each copy-out operation consumes a segment of disk space in the pool.

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. Both files share the same disk space, as shown in the following figure.




Figure 8-1. Example of a New Snapshot's Disk Space Usage

A user then writes to the parent's SMALL.DAT. Before that write happens, Virtual Replicator performs a copy-out to the disk by automatically making a separate physical copy of the *parts* of the file that are being changed. The software then updates the parent disk with the new data.




Figure 8-2. A Copy-out Operation

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 Virtual Replicator software.

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




Figure 8-3. Disk Space Usage After a Copy-Out

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.

Creating Snapshots

You can create a snapshot of either a virtual disk or another snapshot. The disk for which you create the snapshot 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.

NOTE: 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.

To create a snapshot of a disk in a pool:

1. Click **Start > Programs > Compaq SANworks Virtual Replicator > Snapshot Manager**.
2. In the left-hand Scope pane, select the **Snapshot Manager** snap-in for the standalone computer or cluster that the pool is on.
3. Select the virtual disk or snapshot for which you want to create a snapshot, 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. In the Drive Letter Assignment dialog box, choose whether or not to map a drive letter to the new snapshot.

If you do not want to map a drive letter now, accept the default (**No drive letter**), and 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 Snapshot" on Page 8-5.

To map a drive letter, click the **Map drive letter** button, and select a drive letter from the pull-down list. 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.

7. When you assign a drive letter, the new snapshot's volume label is, by default, the same as the parent virtual disk's volume label. This means that multiple disks (a virtual disk and all of its snapshots) have the same volume label. You have the option of changing the name in the Volume Label box. For example, you can change the snapshot's volume label to match its snapshot name. This can be useful for recognizing the snapshot disk in Windows Explorer.

After changing or accepting the volume label, click **Next**.

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

Creating Snapshots in a Cluster

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

Creating Snapshots Automatically

One of the benefits of Virtual Replicator is its ability to automatically create snapshots of virtual disks or other snapshots according to a user-specified schedule. Virtual Replicator provides a set of wizards that lets you schedule the creation of snapshots, as well as other snapshot activities. This functionality is described in Chapter 11, "Automating Tasks."

In addition, Virtual Replicator provides a utility, called SmartSnap, that can automate the number of snapshots maintained for a virtual disk. SmartSnap can be invoked via the command line. The utility is described in Appendix A.

Mapping a Drive Letter to a Snapshot

Use Snapshot Manager, not Disk Administrator, to map drive letters to snapshots. This is the same procedure that you use when mapping a drive letter to a virtual disk.

To map a drive letter to a snapshot:

1. Click **Start > Programs > Compaq SANworks Virtual Replicator > Snapshot Manager**.
2. In the left-hand Scope pane, select the **Snapshot Manager** snap-in for the standalone computer or cluster that the snapshot is on.
3. Select the snapshot, and then right-click **All Tasks > Map Drive Letter**.

4. Select a drive letter from the pull-down list.

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 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.

5. Click **OK**.

The drive letter is mapped to the snapshot, and then the snapshot is available for use by applications such as Windows NT Explorer.

Unmapping a Drive Letter from a Snapshot

Use Snapshot Manager, not Disk Administrator, to unmap drive letters from snapshots. This is the same procedure that you use when unmapping a drive letter from a virtual disk.

To unmap a drive letter from a snapshot:

1. Make sure that no one is accessing the snapshot. You cannot unmap its drive letter if any files on it are open.
2. Click **Start > Programs > Compaq SANworks Virtual Replicator > Snapshot Manager**.
3. In the left-hand Scope pane, select the **Snapshot Manager** snap-in for the standalone computer or cluster that the snapshot is on.
4. Select the snapshot, and then right-click **All Tasks > Unmap Drive Letter**.
5. Click **Yes** to confirm the unmapping.

The drive letter is unmapped from the snapshot, and the snapshot is no longer available to applications.

Using Snapshots to do Backups

The following sections describe how to use snapshots to back up disks. You can use the SnapMgr and NDMgr commands, so that you can create batch jobs to automate your backups. Alternatively, you can use the Snapshot for Backup wizard (see Page 8-10). See Chapter 7, “Managing Virtual Disks,” for information on using snapshots to restore virtual disks.

Backing up a Local Disk Using the Command Line

Follow these steps to use a snapshot to back up a disk:

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. 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 called 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.

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

For example, map drive W to the snapshot:

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

5. 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 2-8).

6. After the snapshot is saved on tape, you can unmap the drive letter from the snapshot and delete the snapshot.

For example:

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

Backing up a Network Disk Using the Command Line

Follow these steps to use a snapshot to do a backup of a disk that is being served as a network disk:

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.

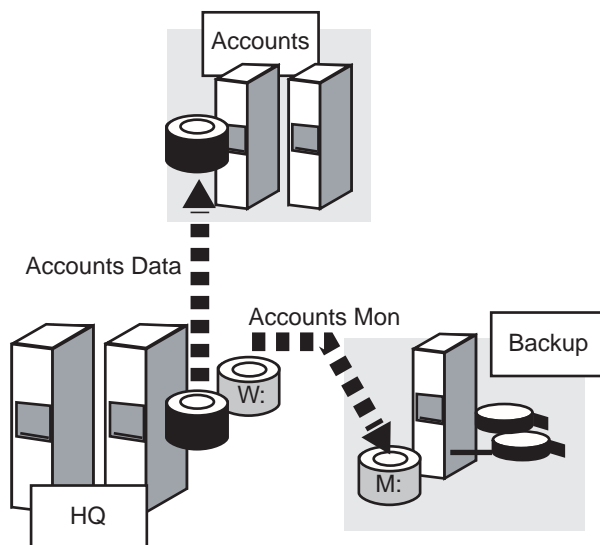


Figure 8-4. Using a Snapshot to Back Up a Network Disk

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. 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
```

8. 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.

NOTE: 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.

9. 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 2-8).

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

For example:

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

11. Stop serving the snapshot.

For example:

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

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

For example:

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

Using the Wizard to Create Snapshots for Backup

Virtual Replicator provides a wizard that automates the procedures described in the “Backing up a Local Disk Using the Command Line” and “Backing up a Network Disk Using the Command Line” sections. The Snapshot for Backup wizard directs you through the creation of a batch script that can include the steps of

- Quiescing an application
- Creating a snapshot
- Restarting the application
- Executing the backup using your backup tool
- Notifying by e-mail or pager when the backup is complete
- Deleting a pre-existing snapshot

See Chapter 11, “Automating Tasks,” for more information on the Snapshot for Backup wizard.

Using Snapshots to Restore Virtual Disks

When you have snapshots of virtual disks, you can use the snapshots to restore individual files or the entire volume in the event of disk corruption. See Chapter 7, “Managing Virtual Disks,” for information on using snapshots to restore virtual disks.

Other Backup and Restore Solutions

The Compaq website contains information on using snapshot technology for backing up and restoring in specific application environments, such as Microsoft Exchange. Go to <http://www.compaq.com/sanworks> to learn more about these solutions.

Displaying Information about Snapshots

To display information about a snapshot, select its family's virtual disk in the left-hand Scope pane of the Snapshot Manager console window. 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	The family to which the snapshot belongs.
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 snapshot, then right-click **Properties**.



The **General** tab shows general information about the snapshot, such as its capacity and disk number. The disk number may change when you restart the computer or when the pool resource fails over within the cluster.

The **Scheduled Tasks** tab lists tasks that have been scheduled for this snapshot. See “Using Windows Scheduled Tasks” on Page 11-3 for more information.

If the snapshot is being served, there are **Network Disk Info**, **Network Disk Statistics**, and **Security** 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 9-17.

Deleting Snapshots

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.

NOTE: You cannot delete a snapshot if you are serving it as a network disk.

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. Click **Start > Programs > Compaq SANworks Virtual Replicator > Snapshot Manager**.
4. In the left-hand Scope pane, select the **Snapshot Manager** snap-in for the standalone computer or cluster where the snapshot is.
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.

NOTE: When you delete a snapshot, make sure to delete any scheduled tasks for the snapshot. Otherwise, if you reuse the old snapshot name for another snapshot, the old tasks will run on the new snapshot. Use the Windows Scheduled Tasks applet to delete scheduled tasks.

Additional methods for deleting snapshots are discussed in Chapter 11, "Automating Tasks" and Appendix E, "Appendix E.."

Chapter 9

Managing Network Disks

This chapter describes how to serve and manage network disks using Virtual Replicator.

The following topics are described in this chapter:

- Serving a Disk (Page 9-4)
- Controlling Which Clients Connect to the Disks (Page 9-8)
- Connecting to a Served Network Disk (Page 9-9)
- Changing the IP Subnet Configuration (Page 9-13)
- Changing the Client System for a Served Network Disk (Page 9-14)
- Using a Served Network Disk as the Cluster Quorum Disk (Page 9-15)
- Tuning Tips (Page 9-16)
- Showing Information about Served Network Disks (Page 9-17)
- Showing Information about Connected Network Disks (Page 9-20)
- Disconnecting from a Served Network Disk (Page 9-21)
- Stopping the Serving of a Network Disk (Page 9-22)

NOTE: 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.

Introduction to Network Disks

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

You can make the virtual disks and snapshots you create available over a network in the following ways:

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

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.

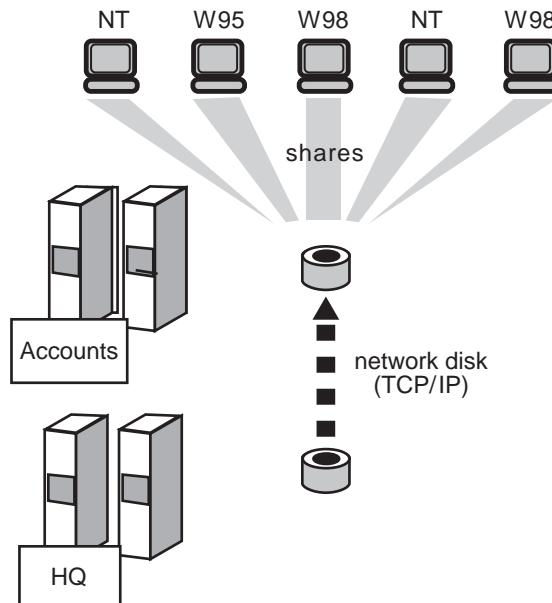


Figure 9-1. Serving a Disk Over a Network

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, many computers can connect to the share. With a network disk, only one standalone computer or cluster can connect directly 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 create shares on it, and some applications cannot install directly onto it or 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. However, you cannot serve a connected network disk.

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 network disk being served has a list of registered clients . This is a list of all the computers that can connect to the served disk. You can 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 for 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.

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 required to serve a disk depend on which type of disk you are serving:

- Serving Virtual Disks and Snapshots (Page 9-4)
- Serving other Disks (Page 9-5)

NOTE: In a cluster, do not use Cluster Administrator to create cluster resources of type SCE Served Disk. Use Network Disk Manager, as described in "Serving other Disks."

Serving Virtual Disks and Snapshots

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

1. Before you serve a virtual disk or snapshot, make sure that:
 - The virtual disk or snapshot has a drive letter mapped to it.
 - All the files on the virtual disk or snapshot are closed. The disk is taken offline when you serve it.
2. Click **Start > Programs > Compaq SANworks Virtual Replicator > Snapshot Manager**.

NOTE: 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.
3. In the left-hand Scope pane, select the virtual disk or snapshot you want to serve, and then right-click **All Tasks**, then **Serve as Network Disk**.
4. You see the welcome screen of the New Network Disk wizard. Click **Next** to start the wizard.
5. 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.

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

This list specifies the clients that can connect to the new disk you are serving. When a remote computer tries to connect to the disk, the request is refused if the name of the computer is not on the list. By default, the list is empty, which means that all computers can connect to the network disk.

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

NOTE: 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.

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

Serving Other Disks 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.

Serving other Disks

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

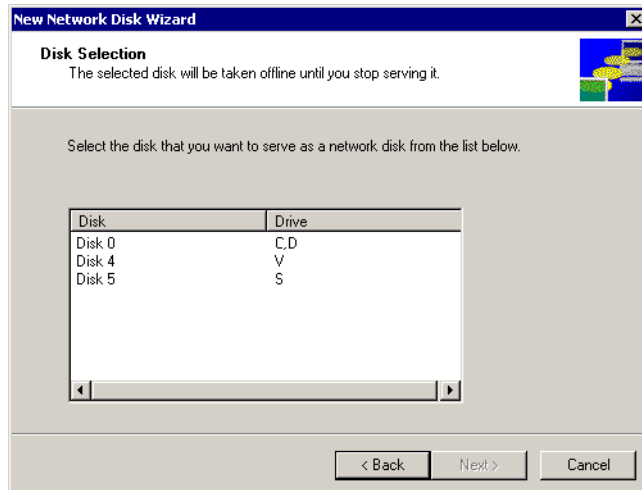
1. Make sure that all the files on the disk are closed. The disk is marked as offline when you serve it.
2. If the disk you want to serve is on a standalone computer, go to Step 4.
3. Make sure that the disk is not configured as a cluster resource of type Physical Disk.

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 C-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. Click **Start > Programs > Compaq SANworks Virtual Replicator > Network Disk Manager**.
9. In the left-hand Scope pane, select the **Served Network Disks** folder, and then right-click **New** then **Network Disk**.
10. You see the welcome screen of the New Network Disk wizard. Click **Next** to start the wizard.
11. The wizard displays a list of disks. Select the one you want to serve, then click **Next**.



NOTE: 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 0 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.

NOTE: Select a disk that has just one drive letter in the **Drive** column. *Compaq SANworks* Virtual Replicator does not support serving a disk with 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.

12. 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.

NOTE: 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.

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

This list specifies the clients that can connect to the new disk you are serving. When a remote computer tries to connect to the disk, the request is refused if the name of the computer is not on the list. By default, the list is empty, which means that all computers can connect to the network disk.

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

NOTE: 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.

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

15. 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.

Serving a Virtual Disk or Snapshot 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.

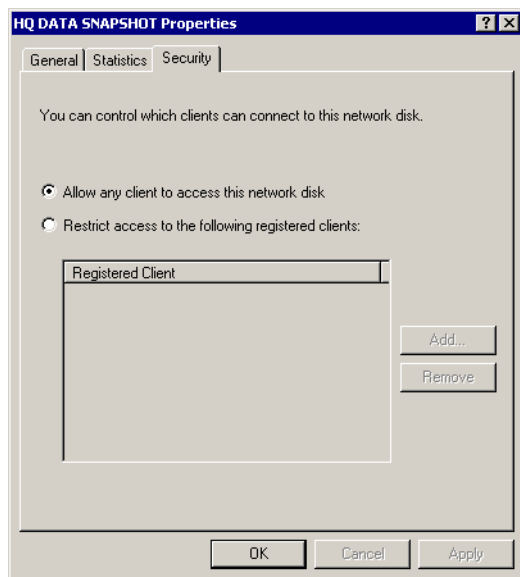
Controlling Which Clients Connect to the Disks

Each network disk that is served by a computer or cluster has a list of the client computers that are registered to connect to it.

Any computer specified on the list can connect to the network disk.

To change the list of registered clients for a network disk:

1. Click **Start > Programs > Compaq SANworks Virtual Replicator > Network Disk Manager**.
2. In the right-hand Scope pane, select the network disk, then right-click **Properties**.
3. Click the Security tab. The Security dialog box is displayed.



Select the **Allow any** button if you want any computer to be able to connect to this network disk.

To control access to the network disk, 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.

NOTE: 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.

4. Click **OK** to save the changes you made.

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 you disconnect it, using either the NDMgr **USE** command (Page C-22) 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.

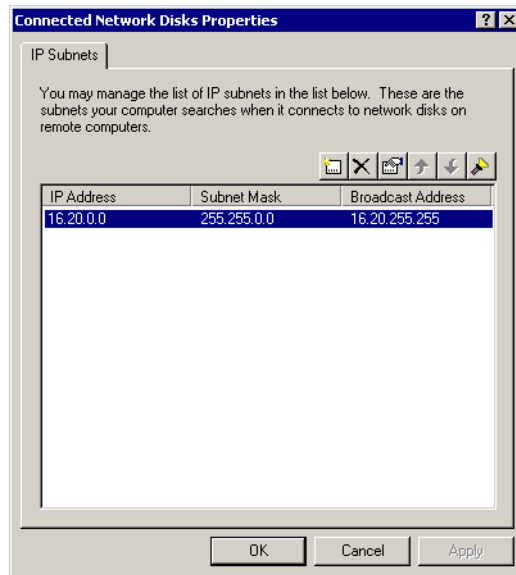
NOTE: In a cluster, do not use Cluster Administrator to create resources of type SCE Connected Disk. Use Network Disk Manager, as described in this section.

To connect to a served network disk:

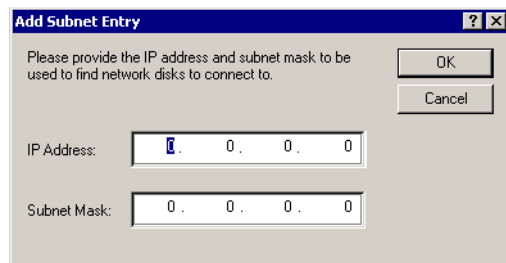
1. Click **Start > Programs > Compaq SANworks Virtual Replicator > Network Disk Manager**.
2. Go to Step 6 if the computer or cluster that you want to connect to the network disk is either in the same IP subnet as the one serving the network disk, or if it is already configured to connect to network disks in that subnet.

Otherwise follow Steps 3 to 5 to configure it to connect to network disks in that IP subnet.

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



4. Click the **New Subnet** toolbar button. You see the Add Subnet Entry dialog box.



- Determine whether or not your gateways are configured to broadcast messages to other IP subnets.

If:	Then:
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 .

- In the left-hand Scope pane, select the **Connected Network Disks** folder, and then right-click **New** then **Connected Network Disk**.
- You see the welcome screen of the Network Disk Connection wizard. Click **Next** to start the wizard.
- 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.
- 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**.
- Click **Finish** to complete the wizard.

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

On Windows NT 4.0, you can use Disk Administrator to map a persistent drive letter. The following steps describe how to do this.

1. 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 C-2) to find out which drive letters are available on every node in the cluster. Choose one of these drive letters.

2. 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 1.

Now assign the drive letter you chose in Step 1 to the network disk. The drive letter is now persistent on the local computer.

3. In a cluster, repeat Step 2 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.

NOTE: You cannot make drive letters persistent on Windows 2000.

Connecting to a Served Network Disk 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.

Changing Settings

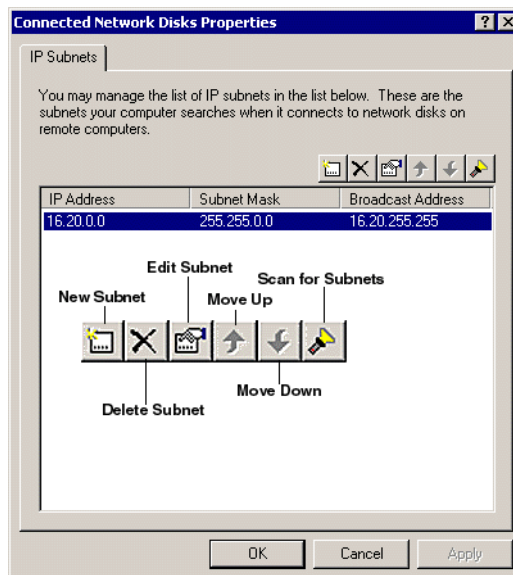
The following sections describe how to change the IP subnet configuration and how to change the client system for a served network disk.

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 present in 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, use the Network Disk Manager snap-in. Select the **Connected Network Disks** folder, then right-click **Properties**.

You see the IP Subnets screen.



You might 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 present in 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.
- If you want 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.

NOTE: 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.

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 attempt to connect the new client system.

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 lets you 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 computer or cluster that is served over the network.

NOTE: 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.

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 9-9.
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 C-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.

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 9-15.

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

Tuning Tips

Use the following suggestions 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, click the **Network Disks** folder in the Network Disk Manager console window, then right-click **Properties**. 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.

Click the **Network Disks** folder in the Network Disk Manager console window, then right-click **Properties**. 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.

Showing Information about Network Disks

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

Showing Information about Served Network Disks

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

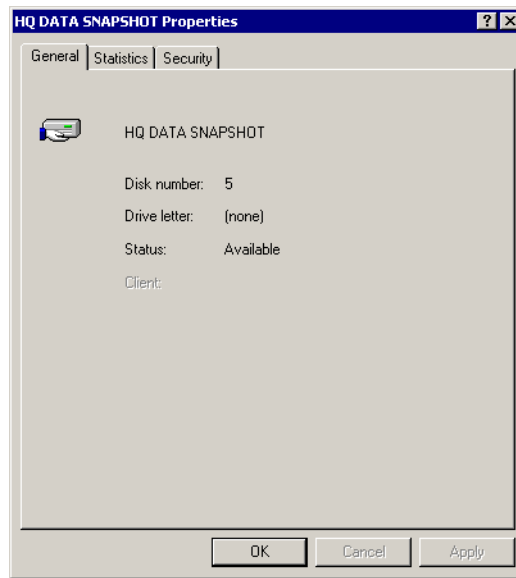
1. Click **Start > Programs > Compaq SANworks Virtual Replicator > Network Disk Manager**.

2. 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.

3. For more information, select the disk then right-click **Properties**. The Properties page contains two tabs: **General** and **Statistics**.



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.

The **Security** tab lets you change the registered clients (see “Controlling Which Clients Connect to the Disks” on Page 9-8 for more information).

Showing Information about Served Network Disks 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.

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. Click **Start > Programs > Compaq SANworks Virtual Replicator > Network Disk Manager**.
2. 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> \netdiskname, 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.

3. 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.

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 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.

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. Click **Start > Programs > Compaq SANworks Virtual Replicator > Network Disk Manager**.
3. In the left-hand Scope pane, select the **Connected Network Disks** folder.
4. In the right-hand Results pane, select the network disk you want to disconnect, then right-click **Delete**.
5. Click **Yes** to confirm that you want to delete the connection.

Disconnecting from a Served Network Disk in a Cluster

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.

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 Explorer.

NOTE: 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 **All Tasks** then **Stop serving**.

NOTE: You cannot stop serving a network disk if any node in the cluster is down.

To stop serving a network disk:

1. Click **Start > Programs > Compaq SANworks Virtual Replicator > Network Disk Manager**.
2. In the left-hand Scope pane, select the **Served Network Disks** folder.

3. 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 9-21.

NOTE: 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.

4. In the right-hand Results pane, select the network disk you want to stop serving, and then right-click **All Tasks** then **Stop serving**.
5. Click **Yes** to confirm that you want to stop serving the disk.

Stopping the Serving of a Network Disk in a Cluster

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.

Chapter 10

Growing Storage Volumes Online

Virtual Replicator provides an online volume growth feature that lets you increase storage capacity without having to restart the computer. With Online Volume Growth, you can grow basic (physical) disks and virtual disks. This feature is available only on Windows 2000.

This chapter describes the following topics:

- Online Volume Growth for Basic and Virtual Disks (Page 10-2)
 - Understanding Basic, Virtual, and Dynamic Disks (Page 10-2)
 - Preparing for Online Volume Growth of Basic or Virtual Disks (Page 10-4)
 - Adding Storage Units to Grow a Pool (Page 10-5)
- Growing Virtual Disks (Page 10-5)
- Growing Basic Disks (Page 10-8)
- Using Windows Explorer to Grow Basic and Virtual Disks (Page 10-22)

Online Volume Growth for Basic and Virtual Disks

The Compaq SANworks Virtual Replicator software enables two different types of online volume growth. You can grow a volume on a basic (physical) disk by creating additional space on the disk into which the volume can expand. You can also grow a virtual disk by adding storage units to the pool and then increasing the virtual disk's size. Virtual Replicator's Online Volume Growth feature is not supported on Windows 2000 dynamic disks, as explained in the following section.

Understanding Basic, Virtual, and Dynamic Disks

The Online Volume Growth feature operates only on basic and virtual disks formatted with NTFS. It is helpful to know the differences between basic and virtual disks, and dynamic disks.

Basic Disks

A basic disk is a standard partitioned disk consisting of a partition table and up to four primary partitions. Configurations of greater than four partitions are possible by creating an extended partition entry (reducing the possible number of primary partitions to three) and creating one or more logical drives. Each primary partition or logical drive is capable of being formatted with a file system and appearing as a volume in the operating system. The table below shows the configuration for a typical basic disk.

Typical Basic Disk Configuration		
Disk 0		
Basic 4.00 GB Online	DIAGS 39 MB FAT Healthy (EISA Configuration)	Windows 2000 (C:) 4.00 GB NTFS Healthy (System)

One of the benefits of using basic disks for online volume growth is simplicity. Even after performing multiple online volume growth operations, the end result is a single partition that has grown over time. The table below shows the configuration for a basic disk after online volume growth.

Basic Disk Configuration After Multiple OVG Operations		
Disk 0		
Basic 192.00 GB Online	DIAGS 39 MB FAT Healthy (EISA Configuration)	Windows 2000 (C:) 192.00 GB NTFS Healthy (System)

Since partitions are recognized by a wide variety of operating systems including DOS, Windows 3.1, Windows 95/98, and Windows NT, and have been used extensively in the past, they must be used if backwards compatibility is a concern. For example, if you dual-boot a system between Windows 95/98 and Windows 2000, you must use basic disks to hold volumes needed under both environments.

Basic disks do not support volume configurations that span across different disk controllers. Previously, this functionality was provided by software fault-tolerance mechanisms such as Volume Sets (RAID0), Stripe Sets (RAID0), Mirror Sets (RAID1), and Stripe Sets with Parity (RAID5). These configurations are not compatible with online volume growth. For more information, see the Windows 2000 documentation about basic disks and basic volumes.

Virtual Disks

Virtual Replicator enables the grouping of basic disks into a concatenated pool of disk space. Virtual disks can then be created from the storage pools. Each virtual disk contains a single partition formatted with the NTFS file system.

Pools created with Virtual Replicator have always been capable of expansion through the addition of new storage units. With the new feature of Online Volume Growth, it is now possible for virtual disks to expand, without rebooting, after new storage units have been added to the pool. For example, you can add physical disks to a RAID set, assign those new storage units to an existing pool, and then use the additional capacity to expand any of the virtual disks associated with the pool.

Dynamic Disks

Dynamic disks are new to Windows 2000 and offer software fault-tolerance features similar to those previously found under Windows NT. Instead of a partition table, dynamic disks rely on a proprietary configuration database that resides near the end of the physical hard disk. Because of this new configuration technique, dynamic disks are not compatible with earlier versions of Windows NT.

Dynamic disks do not require Virtual Replicator's online volume growth to increase in size. However, certain benefits are not available with dynamic disks, such as performing logical drive extension with a Compaq RAID controller. Extending a volume on a dynamic disk requires the creation of a new logical drive on an existing controller or a new controller.

For more information about dynamic disks and dynamic volumes, see the Windows 2000 documentation.

Preparing for Online Volume Growth of Basic or Virtual Disks

Performing online volume growth is straightforward. However, before proceeding to grow your volumes, please review this checklist:

- **Make a reliable backup copy.**

If the disk you plan to grow contains any data you want to keep, it is recommended that you back up the data.

- **Plan for new storage.**

Adding capacity may increase the time required to create backups.

When you increase the size of a virtual disk, it consumes more space in the pool. This reduces the amount of free space that can be used for snapshots of virtual disks. Before performing online volume growth, make sure the pools you have created have adequate free space.

When replacing hard drives or using the Compaq Array Configuration Utility to perform Array Expansion or Logical Drive Extension, the controller must finish rebuilding/expanding/extending before online volume growth can be performed.

Using the Array Configuration Utility to expand an array or extend a logical drive takes time. Examine the current controller rebuild/expand priority settings. See the documentation for the Compaq Array Configuration Utility.

■ Verify volume configuration compatibility.

Under Windows 2000, there are two disk-types that are supported: basic and dynamic. Virtual Replicator's online volume growth operates only on basic and virtual disks formatted with NTFS. In addition, only disks containing NTFS volumes that are not participating in a software fault-tolerance configuration are supported.

NOTE: Dynamic Disks, Volume Sets, Mirror Sets, Stripe Sets, and Stripe Sets with Parity are **not** compatible with online volume growth. Also, volumes formatted with **FAT** and **FAT32** are not compatible; only **NTFS** volumes can be grown.

Although it is possible to convert basic disks to dynamic disks while retaining volumes on the disk, it is not possible to revert back to basic disk after converting to dynamic disk without first deleting the volumes on the disk. To take advantage of online volume growth, **do not convert eligible disks from basic to dynamic**. See "Understanding Basic, Virtual, and Dynamic Disks" on Page 10-2 for more information.

If your volumes are currently using the FAT or FAT32 file systems, it is possible to convert them to NTFS by using the Windows 2000 CONVERT utility. See the Windows 2000 documentation for more information.

Adding Storage Units to Grow a Pool

Before you can grow a virtual disk, you must increase the capacity of its pool by adding storage units to it. As explained in Chapter 6, "Managing Pools," a pool can have up to 8 storage units. You can add storage units to a pool at any time, even when users are accessing its virtual disks and snapshots.

When a storage unit is added to a pool, its size within the pool becomes fixed. Therefore, pool capacity does not increase when you extend a basic disk using a method such as the one described in "Creating Adjacent Space" on Page 10-9. The only way to grow a pool is by adding storage units to it. See "Adding Storage Units to a Pool" on Page 6-9 for more information.

Growing Virtual Disks

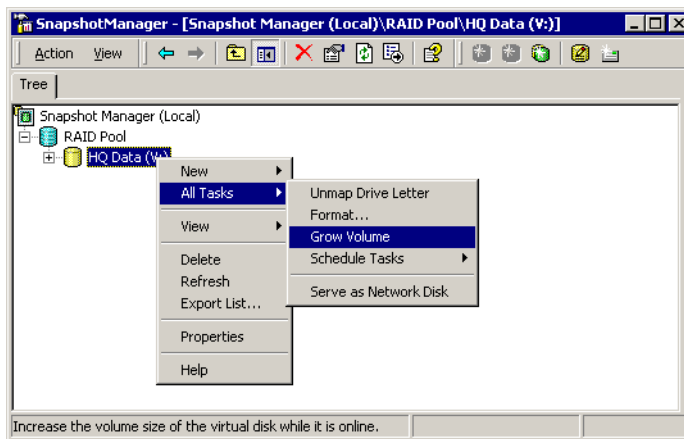
After you add storage units to create pools, you can create virtual disks, as described in Chapter 7, "Managing Virtual Disks." There can be as many as 8 virtual disks in a pool. You can also grow a virtual disk up to a maximum size of 1 terabyte, depending on the pool's properties. Keep in mind, however, that growing a virtual disk depletes free space in the pool that might be needed for snapshots.

The following procedure explains how to grow a virtual disk when using the MMC interface. Online volume growth for virtual disks is also possible using the Online Volume Growth wizard on the Virtual Replicator Start menu. See “Using Virtual Replicator to Grow Basic Disks” on Page 10-10 for this alternate procedure. In addition, this task can be performed by using the command line. See “Grow a Virtual Disk” on Page A-38 for more information.

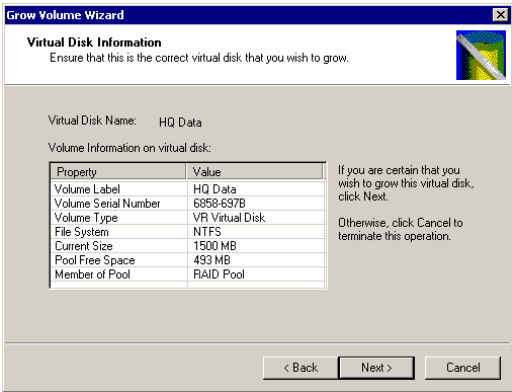
NOTE: Disk volume growth should not be performed at the same time as other major disk management operations, such as defragmenting and disk checking.

To grow a virtual disk:

1. Click **Start > Programs > Compaq SANworks Virtual Replicator > Snapshot Manager**.
2. In the left-hand Scope pane, select the **Snapshot Manager** snap-in for the standalone computer or cluster that the pool is on.
3. Select the virtual disk that you want to grow, and then right-click **All Tasks**, then **Grow Volume**.



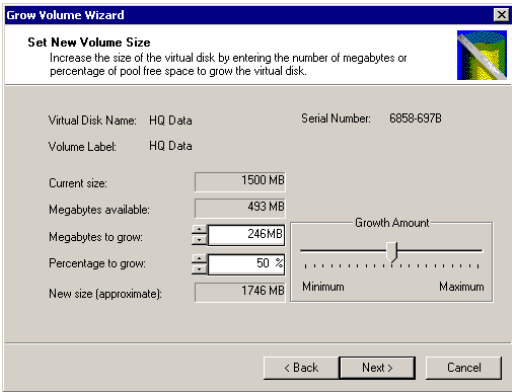
4. You see the welcome screen of the Grow Volume wizard. Click **Next** to start the wizard.
5. In the **Virtual Disk Information** dialog box, confirm that the specified virtual disk is the one you want to grow. If it isn't, click **Back** and repeat Step 4, selecting the correct virtual disk.



Online volume growth operates only on virtual disks formatted with NTFS. If the File System is not NTFS, the wizard presents an error message and instructs you to cancel.

- 6. In the **Set New Volume Size** dialog box, specify the amount by which you want the virtual disk to increase. The **Megabytes available** field shows how much pool free space is available.

Specify the growth amount in the **Megabytes to grow** box, or in the **Percentage to grow** box, or by using the slider control, as shown in the following figure.



Click **Next**.

NOTE: When using Snapshot Manager for online volume growth, the smallest amount a virtual disk can be grown is 1 MB.

7. In the **Confirm Selected Volume** dialog box, confirm online growth for the specified virtual disk. If you want to make changes, click the **Back** button. If you want to proceed with online volume growth, click **Next**.
8. After the virtual disk size increases to its new size, click **Finish** to close the wizard.

Growing Basic Disks

Before you grow a volume on a basic disk, ensure that space is available immediately adjacent to the end of the partition holding the volume as described below.

Preparing Basic Disks for Growth

Only partitions and free space adjacent to the end of a particular volume are eligible to be incorporated into the volume during volume growth, as shown below.

Grow C: into adjacent unused space

Disk 0			
Basic 4.00 GB Online	DIAGS 39 MB FAT Healthy (EISA Configuration)	Windows 2000 (C:) 3.00 GB NTFS Healthy (System)	1.00 GB Unused space

In the diagram below, it is not possible to delete the **DIAGS** volume and grow the **Windows 2000** volume into its place. Only deleting the **OLD TOOLS** volume will supply space for **Windows 2000** to grow into.

Unable to grow C: - no adjacent unused space

Disk 0			
Basic 4.00 GB Online	DIAGS 39 MB FAT Healthy (EISA Configuration)	Windows 2000 (C:) 3.00 GB NTFS Healthy (System)	OLD TOOLS (D:) 1.00 GB FAT

If there is no available adjacent space for a volume to grow into, it may be possible to make space available depending on your particular system configuration, as described in "Creating Adjacent Space."

Creating Adjacent Space

You can make space available for growing volumes in two ways: by deleting a partition and by extending a logical drive.

Creating Adjacent Space Through Deletion

If there is another partition already adjacent to the volume to grow, deleting this partition will make space available. Migrate any information off of the partition to be deleted before actually deleting the partition. In the example below, you can grow the **Windows 2000** volume by deleting the **OLD TOOLS** volume and removing the partition.

Create adjacent space by deleting a volume

Disk 0 Basic 4.00 GB Online	Volume deleted and partition removed		
	DIAGS 39 MB FAT Healthy (EISA Configuration)	Windows 2000 (C:) 3.00 GB NTFS Healthy (System)	1.00 GB Unallocated [OLD TOOLS (D:)]

Creating Adjacent Space Through Logical Drive Extension

You can use the Compaq Array Configuration Utility (ACU) to extend a logical drive on Compaq Array Controllers that support logical drive extension. This method requires a Compaq Array Controller as the underlying disk for your volume. Using a supported controller, a basic disk can be extended using the ACU as shown in the example below. See Compaq Array Configuration Utility documentation for information on extending a logical drive.

1. Configuration before performing logical drive extension

Disk 0 Basic 4.00 GB Online		
	DIAGS 39 MB FAT Healthy (EISA Configuration)	Windows 2000 (C:) 4.00 GB NTFS Healthy (System)

2. Use ACU to extend the basic disk

Disk 0			
Basic 4.00 GB Online	DIAGS 39 MB FAT Healthy (EISA Configuration)	Windows 2000 (C:) 4.00 GB NTFS Healthy (System)	4.00 GB Unused space

3. Grow the volume into the unused space

Disk 0			
Basic 4.00 GB Online	DIAGS 39 MB FAT Healthy (EISA Configuration)	Windows 2000 (C:) 8.00 GB NTFS Healthy (System)	

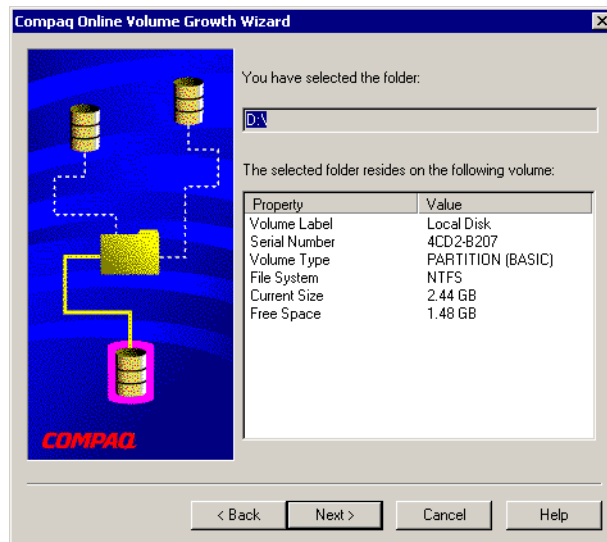
Using Virtual Replicator to Grow Basic Disks

The procedure described below can also be used for growing virtual disks.

NOTE: Disk volume growth should not be performed at the same time as other major disk management operations, such as defragmenting and disk checking.

To grow a basic or virtual disk:

1. Click **Start > Programs > Compaq SANworks Virtual Replicator > Online Volume Growth**.
2. You see the welcome screen of the Grow Volume wizard. Click **Next** to start the wizard.
3. Select a folder that needs additional storage. The folder you select represents a volume that needs to be grown.
4. The next screen shows information about the folder you've selected. Click **Next** to confirm that the specified volume is the one you want to grow. If it isn't, click **Back** and repeat Step 3, selecting the correct volume.

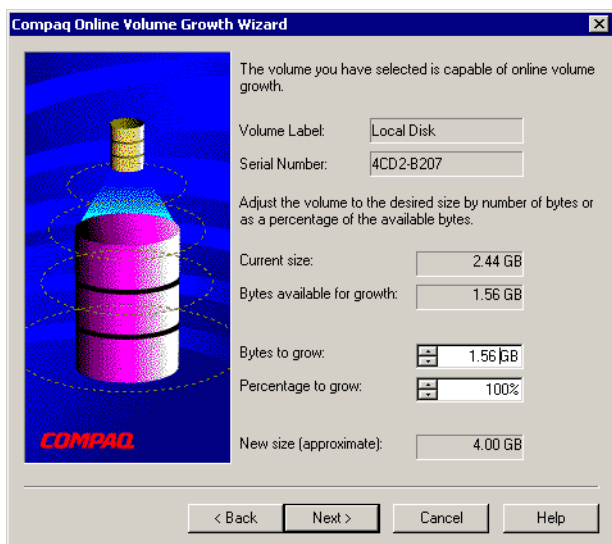


Online volume growth operates only on disks formatted with NTFS. If the file system is not NTFS, the wizard presents an error message and instructs you to cancel the operation.

5. Virtual Replicator determines whether or not the selected folder/volume is capable of online volume growth. If it is, the volume sizing wizard dialog box (see the next figure) shows the current size of the disk and how much capacity is available for growth.

Specify the amount of growth you want in the Bytes to grow box or Percentage to grow box and click **Next**.

NOTE: When using the Online Volume Growth Wizard, the minimum amount a virtual disk can be grown is 1 MB.



6. In the next dialog box, confirm online growth for the specified volume. If you want to make changes, click the **Back** button. If you want to proceed with online volume growth, click **Next**.
7. After the disk size increases to its new size, click **Finish** to close the wizard.

Example of Online Volume Growth for Basic Disks

The following example illustrates how the Online Volume Growth feature can be used to increase storage capacity on a mail and messaging server. The objective is to add storage capacity and make it available to users with little or no interruption in service. The example uses Compaq hardware and software to systematically upgrade the server with minimal impact to server functions.

Configuration

The e-mail server is configured with hard drives connected through an array controller using dual channels. The initial configuration was selected to provide maximum protection against drive, controller, or storage cage failures.

Initial Configuration	
Server	Compaq ProLiant 1600R
Storage Controller	Compaq Smart Array 3200
External Storage Enclosure	StorageWorks Enclosure 4200
Boot Drive (Array A)	2 x 4.3-GB drives (1 internal/1 external) mirrored (RAID1) across two channels
Data Logical Drive (Array B)	3 x 4.3-GB drives (2 internal/1 external) distributed data guarding(RAID5) across two channels
Boot Volume (C:)	A basic disk volume formatted with NTFS. Contains Windows 2000 and any installed application software.
Data Volume (D:)	A basic disk volume formatted with NTFS. Contains application data, e.g., database files.

Initial Storage Configuration														
Channel/ID	0	1	2	3	4	5	6	7	8	9	10	11	12	13
0 (Internal)	4.3 GB	4.3 GB	4.3 GB	none	none	none								Bays: 6 internal/14 external
1 (External)	4.3 GB	4.3 GB	none	none	none	none	none	none	none	none	none	none	none	none

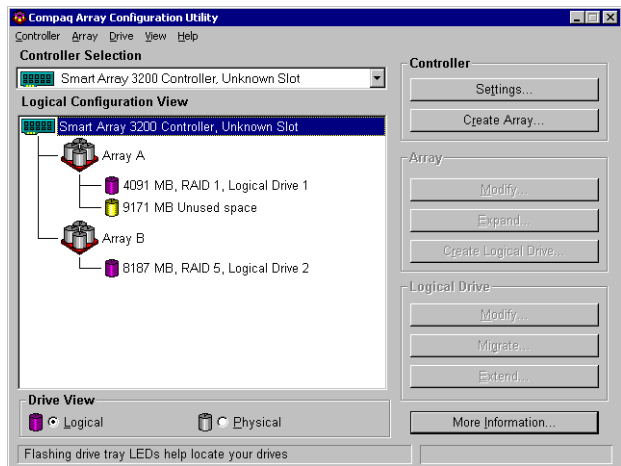
Upgrading and Growing the Boot Drive

After obtaining a reliable backup of the server, the first step is to begin an upgrade of the boot drives comprising the boot volume. Since the logical drive is mirrored, begin by replacing one of the 4.3-GB drives with a 9.1-GB drive and allowing the array controller to rebuild the data on the new drive. Select the drive located at ID 0 in the external enclosure first (see the Initial Storage Configuration table above).

Once the rebuild has finished, repeat the operation with the internal 4.3-GB drive at ID 0. Completion of these steps results in the following configuration.

Storage Configuration after Drive Replacement														
Channel/ID	0	1	2	3	4	5	6	7	8	9	10	11	12	13
0 (Internal)	9.1 GB	4.3 GB	4.3 GB	none	none	none								Bays: 6 internal/14 external
1 (External)	9.1 GB	4.3 GB	none	none	none	none	none	none	none	none	none	none	none	none

Next, use the Compaq Array Configuration Utility to perform logical drive extension. By performing this operation, the existing logical drive will grow to use the new space provided by the 9.1-GB drives.




To access the logical drive extension feature in the Array Configuration Utility:

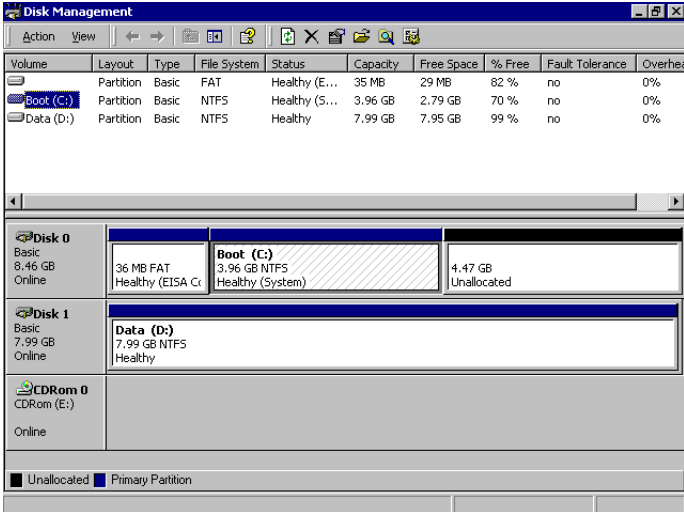
1. Select Array A, **Logical Drive 1**
2. Click the **Drive** menu on the application menu bar
3. Select **Extend Logical Drive...** on the menu.

A dialog box appears to allow adjustment of the current logical drive size.

4. To use all of the available space for the boot drive, drag the slider control completely to the right. This assigns all available space to the logical drive.

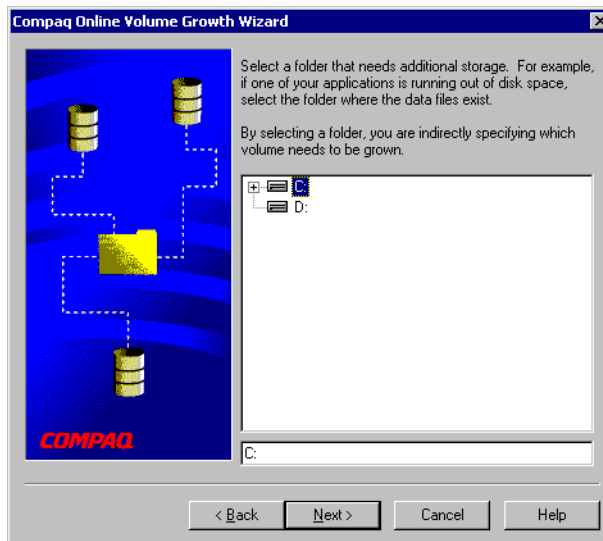


After accepting the changes and saving the new configuration, the array controller begins to move data on the disk. Once it has finished, use **Computer Management > Disk Management** to see the basic disk layout after performing logical drive extension.

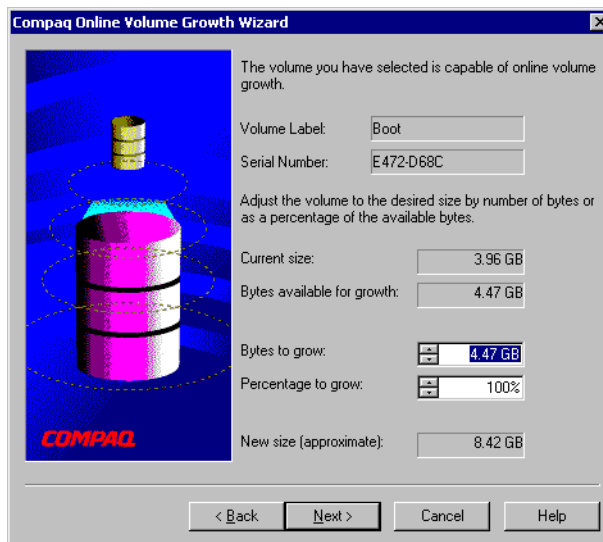


Now use Online Volume Growth to grow the boot volume. Run Online Volume Growth and select the boot volume.

1. Click **Start > Programs > Compaq SANworks Virtual Replicator > Online Volume Growth**.
2. Click **C:**, then click **Next**.

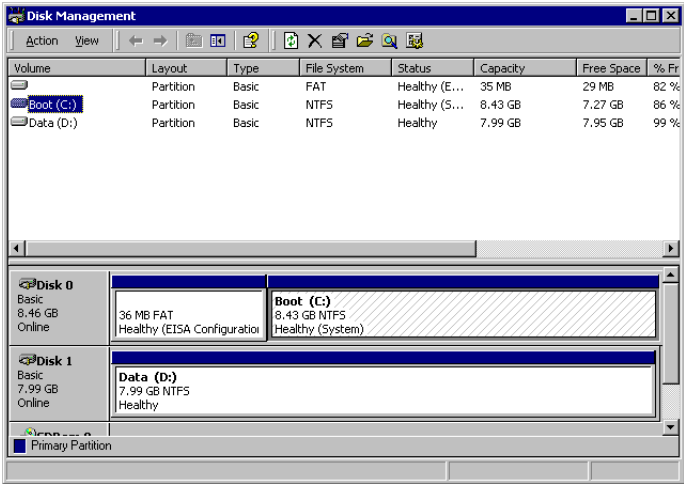


3. In the next dialog box, click **Next** to confirm your selection and continue the wizard.
4. Since all of the available space is to be allocated to the boot volume, select 100% in the volume sizing dialog box, then click **Next**.



5. In the next dialog box, confirm your choices and click **Next**. This completes the Online Volume Growth wizard and creates a larger boot volume without any downtime.

Use **Computer Management > Disk Management** to view the resulting layout of the basic disk.



Growing the Data Drive

This example addresses the need for additional storage on the data drive. The example shows the addition of a single 4.3-GB hard drive; however, it is possible to add as many hard drives as there are available bays in the storage enclosure..

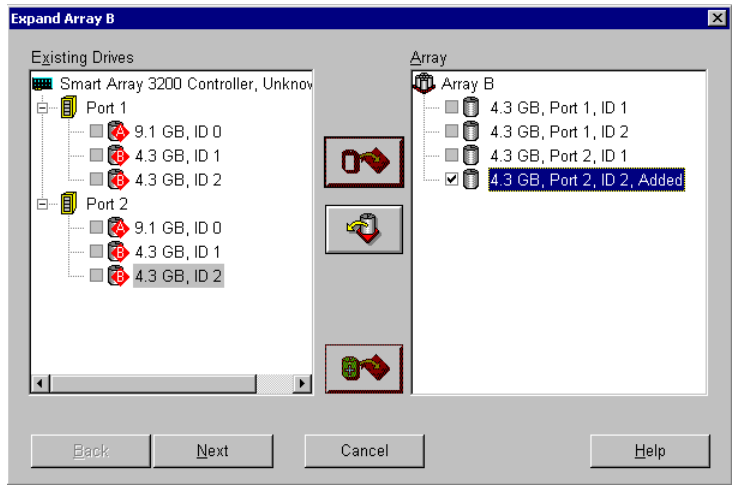
Initial Storage Configuration														
Channel/ID	0	1	2	3	4	5	6	7	8	9	10	11	12	13
0 (Internal)	9.1 GB	4.3 GB	4.3 GB	none	none	none								Bays: 6 internal/14 external
1 (External)	9.1 GB	4.3 GB	none	none	none	none	none	none	none	none	none	none	none	none

To begin the procedure, hot-add the 4.3-GB drive to bay #2 in the external enclosure. The array controller will detect the insertion event and spin-up the drive. Next, run the Array Configuration Utility. Because a new drive is being added to an existing array (Array B), array expansion is necessary. Logical drive extension is also required because the new drive space is being added to the existing logical drive.

System Diagram After Drive Addition

Channel/ID	0	1	2	3	4	5	6	7	8	9	10	11	12	13
0 (Internal)	9.1 GB	4.3 GB	4.3 GB	none	none	none								
1 (External)	9.1 GB	4.3 GB	4.3 GB	none	none	none	none	none	none	none	none	none	none	none

Start by adding the new drive to Array B. Select **Array B** and click the **Expand** button. This will display the **Expand Array B** dialog box. Assign the new 4.3-GB drive on Port 2, ID 2 to the array.



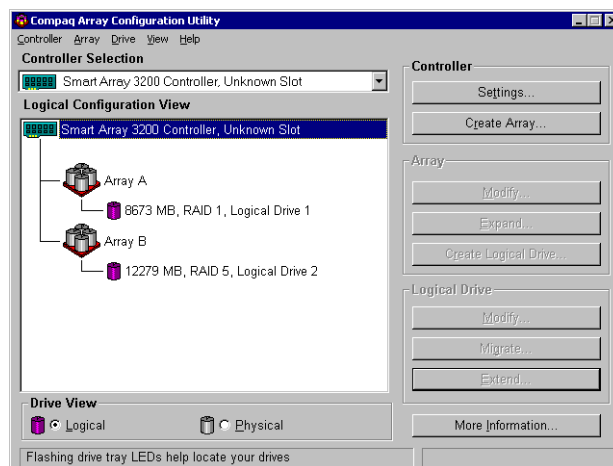
Click the **Next** button.

In the **Create Logical Drives** dialog box, click **Done**. (Creating an additional logical drive is not needed for the purposes of this example.)

Because the logical drive still needs to be extended, it is important to **not save** the configuration at this point. Doing so would require the array controller to finish the array expansion before allowing the logical drive to extend. By waiting to commit the configuration until both operations are prepared, the controller can complete both procedures at one time.

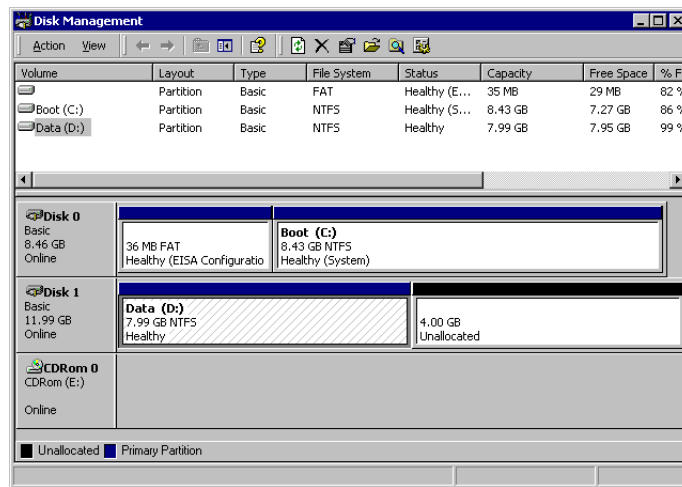
Following the same steps as in “Upgrading and Growing the Boot Drive” on Page 10-13, extend the logical drive in Array B to incorporate the newly added space.

1. Select Array B, **Logical Drive 2**.
2. Click the **Drive** menu, and select **Extend Logical Drive**.
3. In the adjustment dialog box, pull the slider control to the right to allocate all available space.
4. Review the new configuration before saving it to the array controller. It is also useful to review the current rebuild/expand settings on the controller. Click **Settings**. Select **Medium** to allow the controller to balance between servicing regular read/write requests and servicing the expand/extend operation.



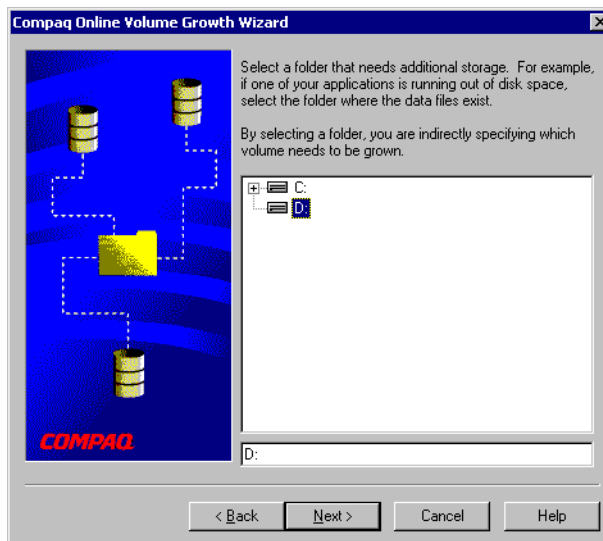
After saving the configuration, the controller will need some time to finish. A good rule of thumb is about 10 minutes/GB, depending on the controller priority settings, array controller load, and hard drive performance.

Use **Computer Management > Disk Management** to see the disk layout after performing logical drive extension.

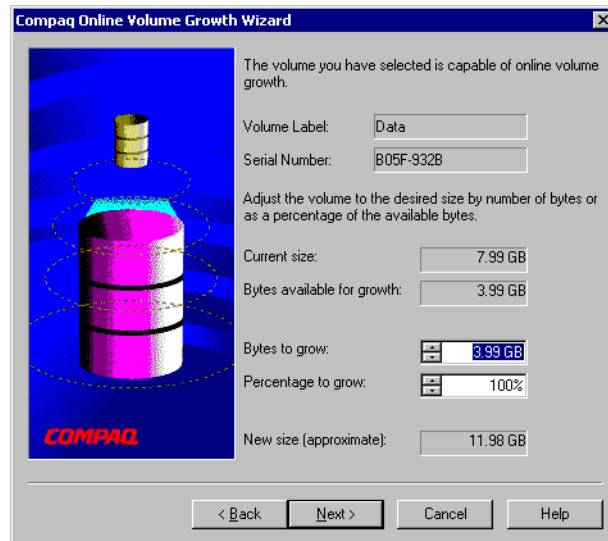


After the controller finishes, the data volume is ready to be grown using Online Volume Growth. To grow the data volume, run Online Volume Growth and select the data volume.

1. Click **Start > Programs > Compaq SANworks Virtual Replicator > Online Volume Growth**.
2. Click **D:**.

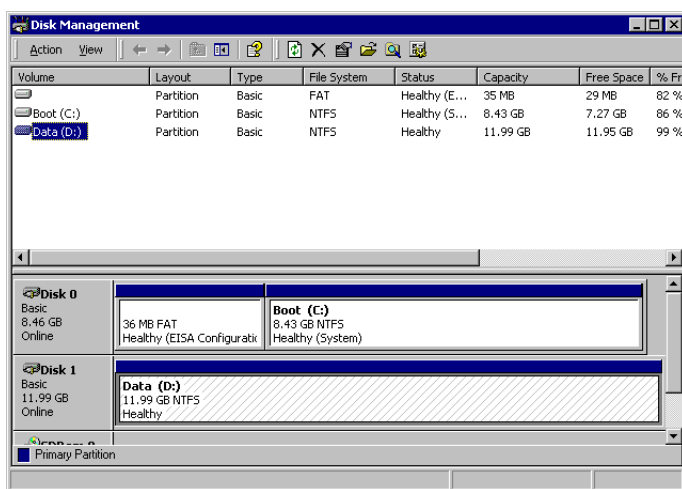


3. Click **Next** in the next dialog box to confirm your selection.
4. Since all of the available space is to be allocated to the data volume, select 100% in the volume sizing dialog box, then click **Next**.



5. Click **Next** in the next dialog box to confirm your selections.
This completes the online volume growth wizard, resulting in a larger data volume without any downtime.

Use **Computer Management > Disk Management** to view the resulting layout of the data disk.



Using Windows Explorer to Grow Basic and Virtual Disks

The Virtual Replicator Volume Growth wizard is also available through the Windows Explorer in two ways:

- As a command on the right-click menu for any virtual or basic disk
- As a tab on the disk's Properties page

To grow a disk using Windows Explorer:

1. Click **Start > Programs > Accessories > Windows Explorer**.
2. Right-click the disk, then click **Volume Growth**, or
Right-click the disk, click **Properties**, click the **Online Volume Growth** tab, then click **Launch**.

The Grow Volume wizard starts.

3. Follow the instructions in "Using Virtual Replicator to Grow Basic Disks" on Page 10-10, starting with Step 2.

Chapter 11

Automating Tasks

Compaq SANworks Virtual Replicator provides Task Automation wizards that allow you to schedule various virtual disk and snapshot activities, such as creating or deleting snapshots, to run automatically. You have access to these wizards after you have created a virtual disk.

The wizards use either the Windows Scheduled Tasks applet or Compaq Batch Scheduler (if it has already been installed) to create schedules for the tasks you specify. You can choose which tool to use from the Virtual Replicator **All Tasks** menu.

NOTE: Using the wizards with Compaq Batch Scheduler requires that you 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 the same system.

This chapter describes the following topics:

- Task Automation Wizards Overview (Page 11-2)
- Setting the Scheduling Method (Page 11-2)
- Scheduling Tasks (Page 11-3 and Page 11-5)
- Viewing and Modifying Task Automation Schedules (Page 11-4 and Page 11-5)

For detailed information on how to use the scheduling wizards, refer to the online help available from the wizards' screens.

Task Automation Wizards Overview

Virtual Replicator provides the following task automation wizards for virtual disks:

- **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.

- **Snapshot For Backup**

The Snapshot For Backup wizard schedules a backup operation of a snapshot to tape.

- **Workday Snapshot**

The Workday Snapshot wizard schedules the creation of a snapshot of a virtual disk at a defined time (for example, start of the work day) and deletion of the snapshot after the defined time (for example, end of the work day).

The following task automation wizards are available for snapshots:

- **Delete Snapshot**

The Delete Snapshot wizard schedules the deletion of a snapshot with a specified name at a predetermined time. If the snapshot name provided does not exist at the time the scheduler runs, no action is taken.

- **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.

- **Snapshot Watchdog**

The Snapshot Watchdog wizard schedules a job, at a specified frequency, that monitors the resource utilization of a snapshot. The snapshot can be optionally deleted if the resource utilization exceeds the present values.

Setting the Scheduling Method

By default, Virtual Replicator performs scheduling tasks by using the built-in Windows Scheduled Tasks applet. If Compaq Batch Scheduler is installed on the system, you can choose either method.

To specify the tool and method for scheduling tasks:

1. Click **Start > Programs > Compaq SANworks Virtual Replicator > Snapshot Manager**.
2. In the left-hand Scope pane, right-click any virtual disk or snapshot.
3. Point to **All Tasks**, and click **Select Scheduler**.
4. In the **Select Scheduling Method** window, click **Compaq Batch Scheduler** or **Windows Task Scheduler**.

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

This sets your preference for all subsequent scheduling activities until you change it again. The following sections describe how to schedule, view, and modify tasks depending on your choice of scheduling tools.

Using Windows Scheduled Tasks

The procedures in this section explain how to use Virtual Replicator for scheduling, viewing, and modifying tasks when you have set the Windows Scheduled Tasks applet as your scheduling tool.

Scheduling Tasks

To schedule a task for a virtual disk or snapshot with the Virtual Replicator wizards:

NOTE: To use the Windows Scheduled Tasks applet, the service needs to be running and you need to have administrative rights.

1. Click **Start > Programs > Compaq SANworks Virtual Replicator > Snapshot Manager**.
2. In the left-hand Scope pane, right-click the appropriate virtual disk or snapshot. Point to the **All Tasks** menu item, and then select **Schedule Tasks**. You see a menu of available wizards.
3. Click your choice of wizards and follow the instructions provided on the wizard screens. See the online help for more detailed information.
4. After you have completed the information required by the wizard, you are then notified that your schedule is about to be entered into the Windows Scheduled Tasks applet. At this point, you have the option of launching the schedule immediately or saving it to launch at another time.

Viewing and Modifying Task Automation Schedules

Tasks scheduled with the Scheduled Tasks applet appear on the Properties page for the virtual disk or snapshot. To view or modify a scheduled task:

1. Click **Start > Programs > Compaq SANworks Virtual Replicator > Snapshot Manager**.
2. In the left-hand Scope pane, right-click the appropriate virtual disk or snapshot. Click the **Properties** menu item
3. On the **Properties** page, click the **Scheduled Tasks** tab. You see a list of tasks and the schedules created with the Windows Scheduled Tasks applet.
4. To modify a schedule, double-click the schedule you wish to modify. This action invokes the Windows properties page for the task.
5. In the properties page, use the Task, Schedule, Settings, and Security tabs to change settings for the task.

You can also use the Scheduled Tasks applet to view and modify scheduled tasks.

With the Windows Scheduled Tasks applet

To view and/or modify tasks using the Scheduled Tasks applet:

1. Double-click **My Computer** on your desktop.
2. On Windows 2000 systems, double-click **Control Panel**.
3. Double-click **Scheduled Tasks**. The window lists all of the tasks.
4. To modify a task, double-click it.
5. In the properties page, use the Task, Schedule, Settings, and Security tabs to change settings for the task.

Using Compaq Batch Scheduler

The procedures in this section explain how to use Virtual Replicator for scheduling tasks when you have set the Compaq Batch Scheduler as your scheduling tool. To view or modify the schedules, use the Compaq Batch Scheduler Coordinator or Monitor.

Scheduling Tasks

To schedule a task for a virtual disk or snapshot with the Virtual Replicator wizards:

NOTE: Before using Compaq Batch Scheduler *to* run a scheduling wizard, you need to log in to the Compaq Batch Scheduler APS. You need a valid Compaq Batch Scheduler user name and password to continue with the wizard.

1. Click **Start > Programs > Compaq SANworks Virtual Replicator > Snapshot Manager**.
2. In the left-hand Scope pane, right-click the appropriate virtual disk or snapshot. Point to the **All Tasks** menu item, and then select **Compaq Batch Wizards**. You see a menu of available wizards.
3. Click your choice of wizards and follow the instructions provided on the wizard screens. See the online help for more detailed information.
4. After 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.

Viewing and Modifying Task Automation Schedules

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

NOTE: Tasks scheduled with the Compaq Batch Scheduler **do not** appear on the Properties page for the virtual disk or snapshot. Use the Compaq Batch Scheduler Coordinator or Monitor to view and modify schedules.

Compaq Batch Scheduler allows you to duplicate a schedule you create for one node to run on another node or a group of nodes.

When you use Compaq Batch Scheduler to duplicate a 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.

Chapter 12

Troubleshooting

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

- Application using a Network Disk Fails or Hangs (Page 12-2)
- Lost Delayed Write Errors (Page 12-4)
- Applications Return Failed Write Errors (Page 12-5)
- Accidentally Delete a Virtual Disk (Page 12-6)
- Snap-in and Commands Do Not Show All the Pools in a Cluster (Page 12-7)
- Disk Structure Corrupt Errors when Accessing a Snapshot (Page 12-7)
- Disk Structure Corrupt Errors when you Stop Serving a Disk (Page 12-8)
- Pool Free Space Fell by More Than the Amount of Data Written (Page 12-9)
- Cannot Connect to a Network Disk in Another Subnet (Page 12-10)
- SCE Connected Disk Resource Fails to Come Online in a Cluster (Page 12-10)
- Reformatting a Virtual Disk Does Not Work (Page 12-11)
- Performance Monitor Shows No Network Disks (Page 12-11)
- Reconstructing a Pool (Page 12-11)
- Tracking Errors (Page 12-12)

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. Click **Start > Programs > Compaq SANworks Virtual Replicator > Network Disk Manager**.

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 9-9.

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 9-9.
 - 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?

<p>Yes</p>	<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 chkdsk on it.</p>
<p>No</p>	<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 9-4.</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.

Lost Delayed Write Errors

NOTE: 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. Click **Start > Programs > Compaq SANworks Virtual Replicator > Snapshot Manager**.
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.

- 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.

- 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.

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

NOTE: 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 6-9.

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.

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:

- Click **Start > Programs > Compaq SANworks Virtual Replicator > Snapshot Manager**.
- 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 6-9.

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.

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.

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.

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.

NOTE: 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 "Backing up a Network Disk Using the Command Line" on Page 8-8.

- 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.

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.

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 defined chunks, called segments. The size of the segment is defined during pool creation.

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 (two 32 KB segments). 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.

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. In the Network Disk Manager console window, select the **Connected Network Disks** folder, then right-click **Properties**.
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.

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.

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.

Performance Monitor Shows No Network Disks

When you serve a network disk, you cannot use the Windows NT Performance Monitor to monitor the served disk. However, clients can monitor any connected network disk.

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” on Page 6-2.
3. Create virtual disks in the new pool. Follow the instructions in “Creating Virtual Disks” on Page 7-2.
4. Use your backup tapes to restore saved data to the new virtual disks.

Tracking Errors

Use the Windows NT Event Log to track errors.

There are three categories of events that produce error, warning, or informational messages:

- SnapDisk driver events
- Network Disk driver events
- Virtual Replicator Lifeguard events

The following sections list the severity, messageId, event source, and event description of each event logged by the Virtual Replicator software.

SnapDisk Driver Events

The **sddriver** messages appear in the system event log.

Severity	MessageId	Event Source	Event Description
Error	2	sddriver	The SdDriver driver failed at initialization.
Error	3	sddriver	Insufficient memory.
Error	4	sddriver	Virtual Replicator cannot create device %2 where %2 is the name of the virtual disk or snapshot that could not be created or initialized.
Error	5	sddriver	Virtual Replicator could not read or write to a disk in the pool.
Error	6	sddriver	Virtual Replicator could not create or initialize disk %2, where %2 is the name of the virtual disk or snapshot that could not be created or initialized.
Error	7	sddriver	Virtual Replicator could not get pointer to device %2, where %2 is the name of the disk device that could not be found.
Error	9	sddriver	Free disk space in the pool %2 (where %2 is the name of the pool) is less than 5 percent. You should delete unneeded snapshots.

Severity	MessageId	Event Source	Event Description
Warning	8	sddriver	Free disk space in the pool %2 (where %2 is the name of the pool) is less than 30 percent. You should delete unneeded snapshots.
Informational	1	sddriver	The SdDriver driver has been started.

Network Disk Driver Events

The **ntds** messages appear in the system event log.

Severity	MessageId	Event Source	Event Description
Warning	1	ntds	Lost connection to \\%2\%3, attempting to reconnect, where %2 is the name of the computer and %3 is the name of the network disk.
Informational	2	ntds	Formed connection to \\%2\%3, where %2 is the name of the computer and %3 is the name of the network disk.

Virtual Replicator Lifeguard Events

The **swvrmon** messages are produced by the Virtual Replicator Lifeguard service, which monitors regular pools and pools with virtual disks containing Microsoft Exchange server data.

Severity	MessageID	Event Source	Event Description
Informational	1	swvrmon	The Virtual Replicator Lifeguard service started.
Informational	2	swvrmon	The Virtual Replicator Lifeguard service stopped.
Error	3	swvrmon	Handler not installed.
Warning	4	swvrmon	Virtual Replicator pool '%2' has reached Microsoft Exchange Server shutdown threshold for partition %3: of < %4 MB free in the pool, where %2 is the pool name, %3 is the drive letter of the virtual disk or snapshot where Exchange has data files, and %4 is the minimum free pool space threshold.
Error	5	swvrmon	Virtual Replicator Lifeguard will now stop Microsoft Exchange Server.
Error	6	swvrmon	Error (%2) stopping service '%3', where %2 is the error number, and %3 is the name of the Microsoft Exchange service.
Informational	7	swvrmon	Virtual Replicator Lifeguard has completed the shutdown of Microsoft Exchange.
Warning	8	swvrmon	Virtual Replicator pool '%2' has reached snapshot deletion policy threshold for partition %3: of < %4 MB free in the pool, where %2 is the pool name, %3 is the drive letter of the virtual disk or snapshot, and %4 is the minimum free pool space threshold (usually 10 MB of disk space).

Severity	MessageId	Event Source	Event Description
Error	9	swvrmon	<p>Virtual Replicator Lifeguard will now execute snapshot delete policy '%2 (%3)' for pool %4 on snapshot %5, where %2 is the type of delete policy, %3 is the delete policy number, %4 is the pool name, and %5 is the snapshot name.</p> <p>The delete policy number (%3) and type (%2) can be any of the following:</p> <ol style="list-style-type: none"> 1 Oldest 2 Oldest with drive letter 3 Oldest without drive letter 4 Newest 5 Newest with drive letter 6 Newest without drive letter 7 Largest delspace 8 Largest delspace with drive letter 9 Largest delspace without drive letter
Error	10	swvrmon	<p>Virtual Replicator Lifeguard will now execute second-chance snapshot delete policy '%2 (%3)' for pool %4 on snapshot %5, where %2 is the type of delete policy, %3 is the delete policy number, %4 is the pool name, and %5 is the snapshot name.</p> <p>The delete policy number (%3) and type (%2) can be any of the following:</p> <ol style="list-style-type: none"> 1 Oldest 2 Oldest with drive letter 3 Oldest without drive letter 4 Newest 5 Newest with drive letter 6 Newest without drive letter 7 Largest delspace 8 Largest delspace with drive letter 9 Largest delspace without drive letter

Appendix 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 4-8.

Command	Description
DRIVES	Shows which drive letters are available 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.

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 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 Commands

```
SnapMgr> DRIVES /NODE:Accounts1
```

```
Available drive letters: GHKMNPQRST
```

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

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
```

Manage a Remote Computer

Syntax

MANAGE *computername*

- *computername* is the name of the standalone computer or cluster that you want to manage.

Use only LANMAN style names, for example, Accounts. DNS style names, such as accounts.domain.com or 123.45.67.122, are not supported.

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

Use this command at the SnapMgr 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 Accounts. Subsequent commands that you enter at the SnapMgr prompt will manage Accounts.

```
SnapMgr> MANAGE Accounts
```

Manage the Local Computer

Syntax

MANAGE /LOCAL

Description

Use this command at the SnapMgr 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
```

Show the Managed Computer

Syntax

MANAGE

Description

Use this command at the SnapMgr 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 Accounts. Subsequent commands issued at the SnapMgr prompt will manage Accounts.

```
SnapMgr> MANAGE  
Managing node Accounts
```

- 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
```

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
```


Create a Pool

Syntax

POOL *pool* /UNITS:*disknumbers* [/SEGSIZE:*segsizesize*]

- *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.

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-30) to find out which storage units are available and what their disk numbers are.
- *segsizesize* is the size in kilobytes of the segment size for the pool. This optional parameter can be one of the following:
32, 64, 128
The default is 32 KB.

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.

A pool can have up to 8 storage units:

- 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-30).
- 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 Commands

```
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
```

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.

There can be 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, then 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: Do 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
```

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 of 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
```

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.

Counter	Description
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
```

SnapMgr Commands

```
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

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*

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.

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. There can be 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
```

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
```


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
```

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
```

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 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 that would free up in the pool if you deleted the snapshot. 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
```

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

VIRTUALDISK

Summary

Use VIRTUALDISK commands to manage virtual disks. You can:

- Create a virtual disk:
VIRTUALDISK *virtualdisk* /POOL:*pool* /CAPACITY:*mbytes*
- Delete a virtual disk:
VIRTUALDISK *virtualdisk* /DELETE
- Format a virtual disk:
VIRTUALDISK *virtualdisk* /FORMAT
- Grow a virtual disk:
VIRTUALDISK *virtualdisk* /GROW:*mbytes*
- Map a drive letter to a virtual disk:
VIRTUALDISK *virtualdisk* /MAP:*drive*
- Show information about virtual disks:
VIRTUALDISK
VIRTUALDISK *virtualdisk*
- Unmap a drive letter from a virtual disk:
VIRTUALDISK *virtualdisk* /UNMAP

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.

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 Virtual Replicator 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-32) 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 Accounts in the News Pool on the local cluster.

```
C:\> SNAPMGR VIRTUAL Accounts /POOL:"News Pool" /CAPACITY:1000
```

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.

NOTE: If you use the /UNMAP switch, the command fails if any files on the virtual disk are open.

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.

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
```

Format a Virtual Disk

Syntax

VIRTUALDISK *virtualdisk* /FORMAT [/ALLOCATIONSIZE:*nsize*]

- *virtualdisk* is the name of the virtual disk that you want to format.
- *nsize* is the size in bytes of the allocation unit. This optional parameter can be one of the following:
512, 1024, 2048, 4096, 8192, 16384, 32768, 65536
If *nsize* is not specified, the operating system sets the allocation unit size (recommended).

Description

This command formats a virtual disk 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.

The optional **/ALLOCATIONSIZE** switch overrides the default allocation unit size, which is set by the operating system.

NOTE: The default setting minimizes the amount of lost space and fragmentation on the volume. It is strongly recommended that you accept the default.

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
```

Grow a Virtual Disk

Syntax

VIRTUALDISK *virtualdisk* /GROW:*mbytes*

- *virtualdisk* is the name of the virtual disk that you want to grow.
- *mbytes* is the capacity in megabytes of the expanded disk. The size is limited by the amount of free space in the pool and the segment size of the pool.

Description

This command grows a virtual disk in a pool.

When you increase the size of a virtual disk, it consumes more space in the pool. This reduces the amount of free space that can be used for snapshots of virtual disks. If there are multiple snapshots in the pool or if snapshot data is rapidly changing, the pool may run out of free space. Therefore, before growing a virtual disk, make sure its pool has adequate free space.

NOTE: Before using the GROW command, always back up the data on the disk.

The command fails if

- You are managing a Virtual Replicator Version 1.1 system or a Virtual Replicator Version 2.0 system on Windows NT 4.0.
- You try to grow a virtual disk over the maximum size of a virtual disk in the pool.
- You try to grow a disk more than the available free space in the pool.

Examples

- 1) This example grows the virtual disk named Accounts Data on the HQ cluster to a new capacity of 2000 megabytes.

```
SnapMgr> MANAGE HQ  
SnapMgr> VIRTUALDISK "Accounts Data" /GROW:2000
```

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
```


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 that would free up in the pool if you deleted the virtualdisk. 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
```

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
```


Appendix B

SmartSnap Command

This appendix describes the SmartSnap utility, which you can use to retain multiple snapshots of a single virtual disk.

SMARTSNAP

Syntax

SMARTSNAP.EXE *virtualdisk* [*UUID*] /KEEP:*n*

- *virtualdisk* is the name of an existing virtual disk of which you wish to take a snapshot
- *UUID* is a unique character string that you specify for Virtual Replicator to identify the task.

This parameter is optional.

- /KEEP: is the keyword to identify the value of snapshot to keep.
- *n* is the number of SmartSnap snapshots you wish to keep.

There can be a total of only 12 snapshots per virtual disk. If the *n* value is not provided the /KEEP: value is set to 12.

Description

SmartSnap is a utility that you can use to retain a specified number of snapshots of a single virtual disk. Normally, when you schedule the creation of a snapshot with Virtual Replicator, the task retains one snapshot for each virtual disk. However, you might want to have a snapshot taken every day and to keep a week's worth of snapshots. Using AT or the Windows Scheduled Tasks applet, you can schedule SMARTSNAP.EXE to do this.

SmartSnap snapshots are identifiable by their snapshot name. The naming convention for a snapshot is:

virtualdisk - truncated to 17 characters if longer

~SCE

a number between 0 and 99

for example, myvdisk~SCE34

Since you might not know the number of snapshots there are of a given virtual disk, SmartSnap examines all the snapshots and determines if there is room for an additional snapshot.

SmartSnap Command

It will delete one of its own snapshots to make room for a new request, always deleting the oldest one first. SmartSnap will only delete snapshots it has created that are not mapped to drive letters. Errors are logged in the application event log.

When the total number of snapshots to be retained has been reached or exceeded, Virtual Replicator deletes the oldest periodic snapshot not mapped to a drive letter. If all snapshots are mapped to drive letters, Virtual Replicator 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 Virtual Replicator limit of 12, and no periodic snapshots can be deleted, then no further snapshots are taken.

Example

- 1) This example shows how to set the number of snapshots to be kept for the virtual disk, myvdisk, to 7.

```
C:\> SMARTSNAP.EXE myvdisk /KEEP:7
```

- 2) To schedule SMARTSNAP.EXE to run on a regular basis, use the Windows Scheduled Tasks applet.

Appendix C

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 4-8.

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.

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
```

NDMgr Commands

- 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
```

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
```

Manage a Remote Computer

Syntax

MANAGE *computername*

- *computername* is the name of the standalone computer or cluster that you want to manage.

Use only LANMAN style names, for example, Accounts. DNS style names, such as accounts.domain.com or 123.45.67.122, are not supported.

Description

Use this command at the 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 Backup. Subsequent commands that you enter at the NDMgr prompt will manage Backup.

```
NDMgr> MANAGE Backup
```

Manage the Local Computer

Syntax

```
MANAGE /LOCAL
```

Description

Use this command at the 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
```

Show the Managed Computer

Syntax

MANAGE

Description

Use this command at the 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 Backup. Subsequent commands issued at the NDMgr prompt will manage Backup.

```
NDMgr> MANAGE
```

```
Managing node BACKUP
```

- 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
```

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 for a served network disk:
REGISTER *netdiskname clientname*
- Remove a registered client for a served network disk:
REGISTER *netdiskname clientname /DELETE*
- Show registered clients for a served network disk:
REGISTER *netdiskname*

Add a Registered Client

Syntax

REGISTER *netdiskname clientname*

- *netdiskname* is the name of the served network disk
- *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 for the specified network disk on the standalone computer or cluster you are managing. Each network disk has a separate list. The list specifies which computers are allowed to connect to the network disk.

By default, the list is empty, which means that all computers can connect to the network disk served by the computer or cluster you are managing.

For each network disk, if the corresponding list is **not** empty, only computers on the list can connect to that disk. When a remote computer tries to connect to the disk, the request is refused if the name of the computer is not on the list.

If you want a cluster to be able to connect to the network disk 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 NetDisk1 on the HQ cluster.

```
NDMgr> MANAGE HQ  
NDMgr> REGISTER NetDisk1 Accounts1
```

- 2) This example adds the computer Accounts2 to the list of registered clients on NetDisk1 on the HQ cluster.

```
C:\> NDMGR HQ REGISTER NetDisk1 Accounts2
```

Remove a Registered Client

Syntax

```
REGISTER netdiskname clientname /DELETE
```

- *netdiskname* is the name of the served network disk
- *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 for the specified network disk on the standalone computer or cluster you are managing. The list specifies which computers are allowed to connect to the specified network disk 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 associated with a network disk, any computer can now connect to the disk.

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 for NetDisk1 on the HQ cluster.

```
NDMgr> MANAGE HQ  
NDMgr> REGISTER NetDisk1 Accounts1 /DELETE
```

- 2) This example removes the computer Accounts2 from the list of registered clients for NetDisk1 on the HQ cluster.

```
C:\> NDMGR HQ REGISTER NetDisk1 Accounts2 /DELETE
```

Show Registered Clients

Syntax

REGISTER *netdiskname*

- *netdiskname* is the name of the served network disk

Description

This command shows the list of registered clients for the specified network disk on the standalone computer or cluster you are managing. The list specifies which computers are allowed to connect to the specified network disk.

Cluster names in the list are ignored.

In a cluster, if you add or remove registered clients when one node in the cluster is down, the lists can be different on different nodes in the cluster.

The command tells you if the lists are different on different nodes in the cluster (see example 3). In this case, use the REGISTER command to add or remove entries to make the lists the same. If the lists are different, clients may not be able to reconnect to network disks when they fail over in the cluster.

Examples

- 1) This example shows the list of registered clients for NetDisk1 on the local cluster. The list is empty, so any computer can connect to the specified network disk.

```
NDMgr> MANAGE /LOCAL
NDMgr> REGISTER NetDisk1
```

```
Clients registered to connect to this network disk:
    ALL
```

- 2) This example shows the list of registered clients for NetDisk1 on the HQ cluster.

```
NDMgr> MANAGE HQ
NDMgr> REGISTER NetDisk1
```

NDMgr Commands

Clients registered to connect to network disk, NetDisk1:

```
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 NetDisk1
```

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 disk, NetDisk1 served by member CENTRAL1:

```
ALL
```

Clients registered to connect to network disk, NetDisk1 served by member CENTRAL2:

```
ACCOUNTS1
ACCOUNTS2
```

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*

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.

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. The disk must have a mapped drive letter before it can be served.

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 C-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:
```

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
```

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 Commands

```
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
```

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*

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 C-22), 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 C-22) show the network disk status as Failed.
- In a cluster, the snap-in and USE command (Page C-22) 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"
```

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
```

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	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). 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.

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:

```
NMGr> MANAGE Accounts
NMGr> 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:

```
NMGr> USE Y:
```

```
Network disk:      \\HQ\ACCOUNTS DATA
Disk number:      7
Status:           Connected
Drive:            Y:
Cluster resource: Yes
Owner:            ACCOUNTS2
```


Appendix D

Virtual Replicator Lifeguard Service

Virtual Replicator Lifeguard is a utility that:

- watches the storage pools for near pool-full conditions and removes snapshots when a pool is nearly full
- watches the storage pools used by Microsoft Exchange 5.5 and shuts down Exchange if the disks used by Exchange are in pools that are nearly full

This appendix describes how the Lifeguard service monitors pools.

See Appendix E for information on setting policies for the Virtual Replicator Lifeguard.

Monitoring Pool Free Space using the Lifeguard Service

The Virtual Replicator Lifeguard service is installed and started automatically. The service checks each drive letter on a system to see if it is a virtual disk or snapshot. If it is, the Lifeguard checks the disk's storage pool for its amount of current free space. This check is performed once every 5 minutes. By default, no snapshots are deleted when approaching pool-full.

You can, however, create policies for the Lifeguard to delete snapshots when pools become full. If the amount of free space goes below a specific threshold, the Lifeguard deletes one or more snapshots to make more space available in the pool. You create these policies by modifying the Windows Registry. See "Setting Policies for the Virtual Replicator Lifeguard" on Page E-6 for more information.

Virtual Replicator Lifeguard posts events in the Application Event log and can be viewed using the Event Viewer. An event will be entered for each disk that exceeds the threshold and additional events will be entered when a snapshot is deleted. See "Virtual Replicator Lifeguard Events" on Page 12-14 for more information.

Monitoring Pool Free Space in a Microsoft Exchange Environment

If you install Microsoft Exchange on a virtual disk, Virtual Replicator Lifeguard monitors the free space in the storage pool that contains the virtual disk, and will shut down Microsoft Exchange in an orderly fashion when a minimum free space threshold is reached.

The Lifeguard determines which disks are being used to store Exchange data. Each disk is then checked to see if it is a virtual disk. If it is a virtual disk, the storage pool that the disk is from is checked for its amount of current free space. If the amount of free space is less than a specific threshold Microsoft Exchange is gracefully shut down.

By default, Virtual Replicator Lifeguard shuts down Microsoft Exchange when the storage pool has only 10MB left. This check is made once every 2 minutes and has very little performance impact upon your system. These defaults can be changed by creating and/or modifying settings in the system registry. See "Setting Policies for the Virtual Replicator Lifeguard" on Page E-6 for more information.

Virtual Replicator Lifeguard posts events in the Application Event log, which can be viewed using the Event Viewer. An event is entered for each disk that exceeds the threshold and additional events are entered when the shutdown of Microsoft Exchange has started and completed.

If there are any errors in shutting down any of the various services used by Microsoft Exchange, an error event will be posted as well. Between the start shutdown and completed shutdown events, various Microsoft Exchange services may post events in the Application and System Event Logs to detail the results of their shutdown.

On a Microsoft Cluster Server, the cluster resources for Microsoft Exchange will be taken off-line. For non-cluster configurations, the local services that make up Microsoft Exchange will be stopped.

Post-shutdown

After Virtual Replicator Lifeguard has shut down Microsoft Exchange, a system administrator can observe the event log to see which disks used by Microsoft Exchange are currently in storage pools that are too full. Using the Virtual Replicator management tools, the system administrator can free up space in the pool by adding disks to the storage pool or deleting other virtual disks or snapshots.

After providing more free space in the storage pool, Microsoft Exchange can be restarted.

Appendix **E**

Virtual Replicator Policy Management

This appendix describes how to set policies to govern the creation and use of Virtual Replicator resources. In many cases, policy management gives you greater control over pools, virtual disks, and snapshots than is generally available through the MMC or command line interfaces.

The rest of this appendix explains how to set the following policies:

- For Pools
 - Maximum number of pools (Page E-3)
 - Minimum amount of free space for a pool (Page E-3)
- For Virtual Disks
 - Maximum number of virtual disks per pool (Page E-4)
 - Allocation unit size for a virtual disk (Page E-4)
- For Snapshots
 - Maximum number of snapshots per pool (Page E-5)
 - Maximum number of snapshots per virtual disk (Page E-5)

In addition, you can create policies that the Virtual Replicator Lifeguard service uses to delete snapshots when pools become full or to shut down Microsoft Exchange. These policies are explained in

- Creating Policies to Delete Snapshots (Page E-7)
- Changing Policies for Microsoft Exchange Pools (Page E-8)

See Appendix D for more information on the Virtual Replicator Lifeguard.

Policy management is handled by editing the Windows Registry key values for Virtual Replicator.

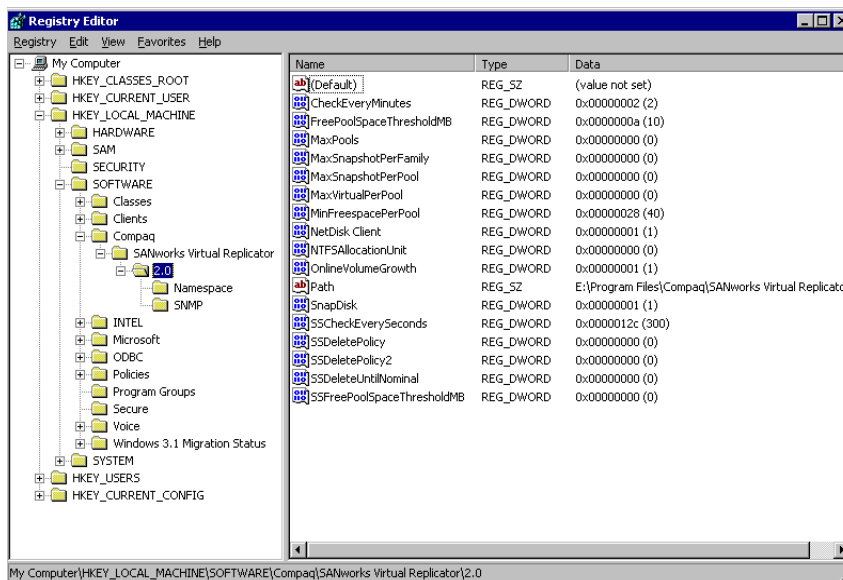


CAUTION: Editing the Windows Registry is advised only for experienced users. If there is an error in your registry, your computer might not function properly. You should be prepared to restore the registry in the event of error. See the Registry Editor Help file for more information.

Using the Registry Editor

All of the procedures for setting policies begin by starting the Registry Editor and opening the Virtual Replicator registry key:

- 1) Click **Start > Run**.
- 2) In the **Run** dialog box, type **regedit** or **regedt32**. The Registry Editor opens.
- 3) In the left pane, open **HKEY_LOCAL_MACHINE\SOFTWARE\Compaq\SANworks Virtual Replicator\2.0**



The right pane of the Registry Editor window shows the value entries for the Virtual Replicator 2.0 key.

Policies can be set by double-clicking the entries listed below.

- CheckEveryMinutes
- FreePoolSpaceThresholdMB
- MaxPools
- MaxSnapshotPerFamily
- MaxSnapshotPerPool
- MaxVirtualPerPool
- MinFreespacePerPool
- NTFSAllocationUnit
- SSDeletePolicy
- SSDeletePolicy2

These are further explained in the sections that follow and in the Virtual Replicator online help.

Setting Policies for Pools

Maximum Number of Pools (MaxPools)

MaxPools sets the maximum number of pools allowed on the system. There is no inherent limit to the number of pools on a system unless this policy is set. Note that the count of pool includes “failed” pools.

Minimum Free Space for a Pool (MinFreespacePerPool)

MinFreespacePerPool reserves a minimum amount of free space in Megabytes that should always be available in the pool. The default setting is 40 MB (decimal). The policy limits virtual disk creation and fails to grow a virtual disk if it would cause the pool free space to fall below the specified threshold.

The 40-MB default ensures that you do not use all available free space in a pool when growing a virtual disk, which could result in lost snapshots and data. It is highly recommended that you do not decrease this minimum amount.

Example 1:

Pool size = 200 MB

MinFreespacePerPool = 40 MB

Creating a virtual disk sized 90 MB would be successful. However, creating a second virtual disk sized 90 MB would fail because a disk that size would leave less than 40 MB of free space. You could create a second virtual disk of up to a maximum of 70 MB (assuming there are no snapshots consuming disk space).

Example 2:

Pool size = 130 MB

MinFreespacePerPool = 40 MB

The pool contains a virtual disk of 60 MB and a snapshot consuming 10 MB (leaving 60 MB of free space).

An attempt to grow the virtual disk by 60 MB fails since this would use all the space in the pool. An attempt to grow the virtual disk by 20 MB succeeds.

Setting Policies for Virtual Disks

Maximum Number of Virtual Disks in a Pool (MaxVirtualPerPool)

MaxVirtualPerPool sets the maximum number of virtual disks per pool. The policy limits virtual disk creation. Virtual Replicator has a hard limit of 8 virtual disks per pool, therefore, only policies of less than 8 are useful.

Virtual Disk Allocation Unit Size (NTFSAllocationUnit)

NTFSAllocationUnit changes the default NTFS allocation unit size for a virtual disk, which is initially set by your operating system. Since, in most cases, the system default provides the highest efficiency for storing data, it is recommended that you not change this value.

If you elect to modify this policy, the revised value becomes the new default. It is used by the Snapshot Manager MMC snap-in when you format a virtual disk and set Allocation Unit to "Default." The value is also used if you do not specify an allocation unit size in the SnapMgr command line interface. You can override this value by picking a nondefault size in the Snapshot Manager console or command line interface or by formatting through Windows Explorer. This policy has NO effect on other Windows NT disks - only on virtual disks formatted through Virtual Replicator tools and commands.

Valid values (in megabytes) for allocation unit size are

- 0 (System default is used)
- 512
- 1024
- 2048
- 4096
- 8192
- 16384
- 32768
- 65536

Setting Policies for Snapshots

Maximum Number of Snapshots in a Pool (**MaxSnapshotPerPool**)

MaxSnapshotPerPool sets the maximum snapshots per pool. This policy limits snapshot creation. Virtual Replicator has a hard limit of 12 snapshots per virtual disk and 8 virtual disks per pool, therefore, only policies of less than 96 are useful.

Maximum Number of Snapshots of a Virtual Disk (**MaxSnapshotPerFamily**)

MaxSnapshotPerFamily sets the maximum number of snapshots per virtual disk. A virtual disk and its child snapshots are known as a family. The policy limits snapshot creation. Virtual Replicator has a hard limit of 12 snapshots per virtual disk (family), therefore, only policies of less than 12 are useful.

Note that the **MaxSnapshotPerFamily** and **MaxSnapshotPerPool** policy are checked simultaneously and whichever limit is reached first will prevent snapshot creation.

Setting Policies for the Virtual Replicator Lifeguard

The Virtual Replicator Lifeguard is a system service that monitors storage pools (see Appendix D). By default, the Lifeguard checks for pool-full situations once every 5 minutes and takes no action. In a Microsoft Exchange environment, the Lifeguard performs a second check of disks every 2 minutes and shuts down Exchange when the storage pool has 10 MB left.

You can change these default behaviors by setting policies using the Window Registry Editor.

If you want to delete snapshots from pools when they become full, you can create a primary deletion policy that specifies

- parameters for the snapshot to be deleted
- the frequency at which the Lifeguard checks the pool
- the minimum threshold for free space
- whether one or multiple snapshots should be deleted at each interval

You can also configure a second-chance deletion policy that executes if the primary deletion policy fails to delete a snapshot. The procedures for creating primary and secondary snapshot deletion policies are explained in “Creating Policies to Delete Snapshots” on Page E-7.

For storage pools in a Microsoft Exchange environment, you can create policies that change

- the frequency at which the Lifeguard checks the Exchange storage pool
- the minimum threshold for free space in the pool

See “Changing Policies for Microsoft Exchange Pools” on Page E-8.

Creating Policies to Delete Snapshots

The following procedure creates a policy whereby the Virtual Replicator Lifeguard deletes one or more snapshots when a pool's free space drops to a specified amount.

To create a delete-snapshot policy:

1. Click **Start > Run**.
2. In the Run dialog box, type **regedit** or **regedt32**. The Registry Editor opens.
3. In the left pane, double-click **HKEY_LOCAL_MACHINE\SOFTWARE\Compaq\SANworks Virtual Replicator\2.0**
4. To create a primary delete-snapshot policy, double-click **SSDeletePolicy** in the right pane. This policy is executed first.
5. In the DWORD Editor box, type one of the following values to specify which snapshot to delete:
 - 0 - None
 - 1 - Oldest snapshot
 - 2 - Oldest snapshot with a drive letter assigned
 - 3 - Oldest snapshot with no drive letter assigned
 - 4 - Newest snapshot
 - 5 - Newest snapshot with a drive letter assigned
 - 6 - Newest snapshot with no drive letter assigned
 - 7 - Snapshot with the largest delspace
 - 8 - Snapshot with the largest delspace with a drive letter assigned
 - 9 - Snapshot with the largest delspace with no drive letter assignedClick **OK** to close the DWORD Editor box.
6. To create a secondary delete-snapshot policy, double-click **SSDeletePolicy2** in the right pane. This policy is executed only if the primary delete policy doesn't execute. For example, if your primary policy specifies to delete the oldest snapshot with a drive letter assigned, your secondary policy could specify the oldest snapshot without an assigned drive letter. The secondary policy would execute only if the primary policy parameters couldn't be met.
7. In the DWORD Editor box, specify a value as in Step 5. Click **OK** to close the DWORD Editor box.

8. By default, the Lifeguard checks regular storage pools at intervals of 5 minutes (300 seconds). To change the number of seconds between checks, double-click **SSCheckEverySeconds** in the right pane.
9. In the DWORD Editor box, specify a value greater than or equal to 10 (seconds). The minimum is 10 seconds. Click **OK** to close the DWORD Editor box.
10. To change the minimum amount of free space that can be in a pool before the snapshot-delete policy is executed, double-click **SSFreePoolSpaceThresholdMB** in the right pane.
11. In the DWORD Editor box, specify a value greater than or equal to 0 (MB). Click **OK** to close the DWORD Editor box.
12. To specify the number of snapshots to delete, double-click **SSDeleteUntilNominal** in the right pane.
13. In the DWORD Editor box, specify a value.
 - A value of 0 results in the deletion of 1 snapshot every interval (SSCheckEverySeconds) from any pool whose free space is less than the defined threshold amount (SSFreePoolSpaceThresholdMB).
 - A value of anything other than 0 (1 or more) results in the deletion of up to 12 snapshots from any pool whose free space is less than the defined threshold amount.Click **OK** to close the DWORD Editor box.
14. Close the Registry Editor after making your changes.

NOTE: Virtual Replicator Lifeguard must be restarted for these new settings to take effect. Use the Control Panel to stop and start the Virtual Replicator Lifeguard service, or type NET STOP SWVRMON followed by NET START SWVRMON from a Command Prompt.

Changing Policies for Microsoft Exchange Pools

The Virtual Replicator Lifeguard checks Microsoft Exchange storage pools and performs an orderly shutdown of Exchange when the pools run out of free space.

To change Lifeguard policies for Exchange storage pools:

1. Click **Start > Run**.
2. In the Run dialog box, type **regedit** or **regedt32**. The Registry Editor opens.
3. In the left pane, double-click **HKEY_LOCAL_MACHINE\SOFTWARE\Compaq\SANworks Virtual Replicator\2.0**

Virtual Replicator Policy Management

4. By default, the Lifeguard checks Microsoft Exchange storage pools every two minutes. To change the number of minutes between checks, double-click **CheckEveryMinutes** for Exchange pools.
5. In the DWORD Editor box, specify a value greater than or equal to 1 (minute).
6. Lifeguard shuts down Exchange when a pool falls below a specified threshold (10 MB by default). To change the minimum amount of free space that can be in a pool, double-click **FreePoolSpaceThresholdMB**.
7. In the DWORD Editor box, specify a value greater than or equal to 0 (MB).
8. Close the Registry Editor after making your changes.

NOTE: Virtual Replicator Lifeguard must be restarted for these new settings to take effect. Use the Control Panel to stop and start the Virtual Replicator Lifeguard service, or type NET STOP SWVRMON followed by NET START SWVRMON from a Command Prompt.

Glossary

allocation unit	The smallest amount of disk space that can be allocated to hold a file. The operating system sets the default allocation unit size for a volume to minimize the amount of lost space and fragmentation.
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.
basic disk	A name given to a standard physical drive that contains a partition table. The partition table holds up to four partitions, however, defining an extended partition allows for more than four partitions by creating logical drives (not to be confused with array controller logical drives).
cluster	Two or more independent computers that are accessed and managed as a single system. <i>See also</i> 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 Virtual Replicator 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	<p>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:</p> <p>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.</p> <p>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.</p> <p>The Delspace value is rounded down to the nearest MB.</p>
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.
dynamic disk	A physical drive that is managed using one or more extents instead of partitions. Extents are regions that correspond to a set of sectors on a particular disk drive. An extent database resides on sectors relative to the end of the drive.
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.
extended partition	A special partition container that allows basic disks to describe more than four partitions (primary or logical drive).
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	<p>A set of disks in a pool that are derived from one another and that can share disk space.</p> <p>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 Virtual Replicator software.</p>

hub	Network hardware that joins communications lines together in a star configuration.
Lifeguard	A Virtual Replicator system service that monitors storage pools.
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 <i>about</i> 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. <i>See also</i> cluster.
names	<i>See</i> 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 Virtual Replicator. To use this file, click Start, then Programs, then Network Disk Manager .

network disk name	<p>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.</p> <p>Do not confuse this name with the volume label that you see in Windows Explorer. When the client connects to the disk, the client sees the volume label in Explorer, not the network disk name.</p> <p>The network disk name is used only by Virtual Replicator 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.</p>
node	<p>A computer that is a member of a cluster.</p>
online backup	<p>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.</p>
parent	<p>When you create a snapshot, the disk that you take the snapshot of is called the parent.</p>
partition	<p>A region defined by a starting offset and a sector count that describes all or part of the usable sectors on a basic disk. A partition resides only on a single basic disk and cannot be combined with other partitions. A single partition forms a volume</p>
pool	<p>A collection of one or more physical disks or controller-based fault-tolerant disk sets that are used to store virtual disks and snapshots.</p>
possible owners	<p>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.</p>
preferred node	<p>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.</p> <p>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.</p>
quorum	<p>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.</p>

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	<i>See</i> Small Computer System Interface.
segment	The smallest unit of disk space that is copied during a copy-out operation. It is a fixed-sized unit of contiguous bytes. The default size of a segment depends on the size of the disks in the pool.
segment size	A value that is used for determining the largest possible virtual disk in a pool. The segment size is set when a pool is created.
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	<p>A disk in a pool that was created by making a virtual copy of another disk, called the parent disk.</p> <p>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.</p> <p>To applications and users, a snapshot looks like a normal disk. You can read and write it like a normal disk.</p> <p>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.</p>

Snapshot Manager	<p>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.</p> <p>A default console file is provided with Virtual Replicator. To use this file, click Start bar, then Programs, then Snapshot Manager.</p>
snapshot name	<p>A name that you supply when you create a snapshot.</p> <p>Do not confuse this name with the volume label that you see in Windows Explorer. The snapshot inherits the same volume label as its parent.</p> <p>The snapshot name is used only by Virtual Replicator software. You cannot change it once you have created the snapshot.</p>
standalone computer	<p>A computer that is not in a cluster.</p>
storage pool	<p><i>See pool.</i></p>
storage unit	<p>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.</p>
virtual disk	<p>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.</p>
virtual disk name	<p>A name that you choose when you create a virtual disk.</p> <p>Do not confuse this name with the volume label that you see in Windows 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 Explorer, the virtual disk name does not change.</p> <p>The virtual disk name is used only by the Virtual Replicator software. You cannot change it once you have created the virtual disk.</p>
volume	<p>An abstraction of partitions and extents that describes a bounded set of available sectors suitable for use by a file system.</p>
volume label	<p>A name that is used to identify a particular volume.</p>

Index

A

- abbreviating commands 4-11
- Action menu 4-4
- adding
 - computers to registered client lists 9-5, 9-7, C-9
 - IP subnets to client configurations 9-13
 - snap-ins to MMC console files 4-3, 4-5
 - snapshots to pools 8-3
 - storage units to pools 6-9, A-12
 - virtual disks to pools 7-2
- adjacent space 10-8 to 10-10
- Adobe Acrobat Reader 3-3
- allocation unit A-37
- allocation unit size 7-3, 7-6, E-4
- Alpha serving network disks to Intel
 - not supported 2-9, 12-8
- application event log 12-2, 12-10
- applications fail or hang 12-2, 12-5
- array expansion 10-18
- automatic snapshots 8-5
- available drive letters A-2, C-2

B

- backup
 - local disk 8-7
 - network disk 8-8
 - online 1-9, Glossary-4
 - Snapshot for Backup wizard 8-10, 11-2
 - solutions 8-11
 - using snapshots for 8-7, 8-10

basic disks

- adjacent space 10-8
 - configuration 10-2
 - differences with dynamic disks 10-2
 - growing 10-10
 - partitions 10-2
 - unsupported volume configurations 10-3
- benefits of Virtual Replicator 1-9
- boot drive
 - upgrading 10-13 to 10-15
- boot volume
 - growing 10-15 to 10-17

C

- cache data 8-3
- capacity
 - increasing 1-5
 - of pools 2-4, 10-5, A-16
 - of snapshots A-21, A-29
 - of storage units A-30
 - of virtual disks 1-2, 1-3, 7-2, A-42
- case insensitivity
 - in commands 4-11
 - in names 2-10
- client software 2-9
- Cluster Administrator
 - when not to use 6-5, 6-11
- cluster groups
 - creating A-10, C-16, C-24
 - defined Glossary-1
 - deleting A-15, C-18, C-27

Compaq SANworks Virtual Replicator System Administrator's Guide

- cluster names 2-11
 - defined Glossary-1
 - cluster resources
 - creating 4-13
 - defined Glossary-1
 - deleting 9-22
 - in a pool 6-5, 6-11
 - Physical Disk resource 9-5, C-16
 - SCE Connected Disk resource 4-13, 12-10, C-24, C-27
 - SCE Pool resource 4-13, 6-5, 12-7, A-10, A-15
 - SCE Served Disk resource 4-13, C-16
 - SCE Storage Unit resource 4-13, 6-5, 6-11, A-15
 - clusters
 - adding storage units to a pool 6-11
 - and network disks 9-9
 - creating pools 6-5
 - defined Glossary-1
 - installing cluster-aware applications 1-2
 - installing Virtual Replicator 3-3
 - mapping drive letters to virtual disks 7-5
 - serving virtual disks 7-2
 - with nodes 100 miles apart 9-15
 - commands
 - abbreviating 4-11
 - how to use 4-8
 - managing remote computers and clusters 4-10
 - summary of 4-9
 - upper or lower case in 4-11
 - privileges required 4-12
 - DRIVES A-2, C-2
 - MANAGE A-4, C-4
 - POOL A-8
 - REGISTER C-8
 - SERVE C-14
 - SMARTSNAP B-2
 - SNAPSHOT A-20
 - USE C-22
 - VIRTUALDISK A-32
 - Compaq
 - general information 3-3
 - Compaq Array Configuration Utility 10-9, 10-14
 - Compaq Array Controllers 10-9
 - Compaq Batch Scheduler 11-5
 - computer names
 - defined Glossary-1
 - computers, standalone
 - defined Glossary-6
 - connecting to network disks 9-9, C-23
 - restrictions 2-9
 - console files
 - adding snap-ins to 4-3, 4-5
 - default Virtual Replicator files 4-3
 - explained 4-3
 - copy-out operations
 - caused by defragmenting disks 7-7
 - defined Glossary-2
 - explained 8-2
 - getting statistics on 6-14, A-17
 - predicting 2-5
 - Create Snapshot wizard 11-2
 - creating
 - cluster groups A-10, C-16, C-24
 - cluster resources 9-1, A-10, C-16, C-24
 - connections to served network disks 9-9, C-23
 - pools 6-2
 - served network disks 9-4, C-15
 - virtual disks 7-2
- D**
- data mining and snapshots 1-10
 - data volume
 - growing 10-17, 10-20
 - defragmenting virtual disks 7-7
 - Delete Snapshot wizard 11-2

- deleting
 - cluster groups A-15, C-18, C-27
 - cluster resources 9-1, A-15, C-27
 - connections to served network disks C-27
 - pools 6-15, A-15
 - pools in a cluster 6-1, 6-16
 - served network disks 9-22, C-18
 - snapshots 8-13, A-26
 - virtual disks 7-10, A-35
- Delspace
 - defined Glossary-2
 - explained 6-8
 - showing 6-8, A-29, A-42
- disconnecting from network disks 9-21, C-27
- disconnecting from network disks in a cluster 9-22
- disk space
 - available in a pool 6-5, 6-8, 12-4, 12-5
 - used by a snapshot 2-4, 6-9, 7-7, 8-2, 8-13, A-26, A-29
 - used by a virtual disk 2-4, A-42
- disks
 - see network disks
 - see pools
 - see snapshots
 - see storage units
 - see virtual disks
- documentation 3-3
- drive letters
 - mapping to snapshots 8-4, 8-5, A-23
 - mapping to virtual disks 7-1, 7-2, 7-4
 - showing available A-2, C-2
 - unmapping from snapshots 8-6, A-25
 - unmapping from virtual disks 7-5, A-43
- DRIVES command A-2, C-2
- dynamic disks
 - configuration 10-4
 - differences with basic disks 10-4
 - extending volumes 10-4
 - fault tolerance 10-4
 - logical drive configuration 10-4
- E**
 - error messages
 - failed to bring SCE Connected Disk resource online 12-2, 12-10
 - failed to bring SCE Pool resource online 12-7
 - failed write 12-5
 - lost connection to network disk C-25
 - lost delayed write 12-4
 - pool 70% full 6-6
 - pool 95% full 6-6
 - event log 6-6, 12-2, 12-7, 12-10, C-25
 - Exchange 1-6
 - extension snap-ins for MMC 4-2
- F**
 - family
 - explained 8-2, A-33
 - showing 7-8, 8-11, A-28, A-41
 - FAT file system
 - not supported 2-3, 2-10
 - fault tolerance
 - dynamic disks 10-4
 - software, not supported 6-3
 - supported 6-2
 - formatting virtual disks 7-3, 7-6, 12-11, A-37
 - free space in pools
 - and Delspace 6-8
 - increasing 8-13
 - monitoring 6-5
 - specifying E-3
- G**
 - glossary of terms Glossary-1
 - growing
 - basic disks 10-8 to 10-12
 - virtual disks 10-5 to 10-8, A-38
- I**
 - I/O statistics
 - for network disks 9-19, C-20
 - for pools 6-14, A-17
 - I/Os to a network disk stall 12-2, C-25

installation

- components 3-2
- in a cluster 3-3
- Internet Explorer requirements 3-1
- operating system requirements 3-1
- prerequisites 3-1
- procedure 3-2

Intel serving network disks to Alpha
not supported 2-9, 12-8

Internet Explorer
installation requirement 3-1

IP subnets
configuring network disk clients 9-13

L

Lifeguard service D-1 to D-3
policies E-6

list of registered clients
see registered clients

local disks
defined Glossary-3

log, error 6-6, 12-2, 12-7, 12-10, C-25

logical drives
dynamic disks 10-4
extending 10-9, 10-14, 10-19

M

MANAGE command A-4, C-4

Management Information Base 1-5

mapping drive letters
in a cluster 7-5
to a virtual disk 7-4
to snapshots 8-4, 8-5, A-23
to virtual disks 7-1, 7-2, 7-4

MIB 1-5

Microsoft applications 1-6

Microsoft Exchange
monitoring pool free space D-2
setting policies for Exchange pools E-8

Microsoft Management Console

see MMC

mirror sets, host-based
not supported 2-3, 2-10

MMC

Action menu 4-4
adding snap-ins 4-5
console files 4-3
defined Glossary-3
Results pane 4-3
right-click menus 4-4
Scope pane 4-3
understanding 4-2
View menu 4-4

monitoring

I/Os to network disks 9-19, C-20

I/Os to pools 6-14, A-17

pool free space 6-5, D-2

monitoring performance 1-6

.msc files 4-3

MSCS cluster

defined Glossary-3

multiple partitions
not supported 2-3, 2-10

My Computer 7-9

N

names

case insensitivity of 2-10
cluster names 2-11, Glossary-1
computer names Glossary-1
conventions 2-10
network disk names Glossary-3
snapshot names Glossary-5
uniqueness 2-11
virtual disk names 7-2, Glossary-6

NDMgr commands
 abbreviating 4-11
 how to use 4-8
 managing remote computers and clusters 4-10
 privileges required 4-12
 summary of C-1
 upper or lower case 4-11
 DRIVES C-2
 MANAGE C-4
 REGISTER C-8
 SERVE C-14
 USE C-22

Network Disk Manager extension snap-in 4-2

Network Disk Manager snap-in 4-2

network disks
 backup 8-8
 benefits of 1-9
 comparison with shares 9-3
 connecting to 9-9, C-23
 defined Glossary-3
 disconnecting from 9-21, C-27
 disconnecting in a cluster 9-22
 explained 1-4, 9-2
 in a cluster 9-9
 names of Glossary-3
 performance of 9-3, 9-16
 properties 9-17
 restrictions on use of 2-10
 serving 9-4, C-15
 showing information about 9-17, 9-20, C-19, C-29
 stopping serving 9-22, C-18
 supported configurations 2-9
 tuning tips 9-16
 using as cluster quorum resource 9-15, C-25

node
 defined Glossary-4

O

online backup 1-9, 8-7, Glossary-4

online volume growth 1-5
 array expansion 10-18
 basic disks 10-2, 10-10
 boot volume, growing 10-15 to 10-17
 Compaq Array Configuration Utility 10-9
 creating adjacent space 10-9
 data volume 10-17, 10-20
 dynamic disks 10-4
 example 10-12 to 10-22
 extending logical drives 10-19
 partitions 10-3
 preparing for 10-8
 storage controllers 10-19
 unsupported configurations 10-3
 using Windows Explorer 10-22

Oracle 1-6

P

parent
 defined Glossary-4

partitioning virtual disks 7-6

partitions
 backwards compatibility 10-3
 basic disks 10-2
 deleting to create space 10-9
 online volume growth 10-3

partitions, FAT
 not supported 2-3

partitions, multiple
 not supported 2-3, 2-10

Performance Monitor 1-6

performance of network disks 9-3, 9-16

Physical Disk cluster resources 9-5, C-16

planning Virtual Replicator configurations 2-1

policies
 for deleting snapshots E-6, E-7
 for pools E-3
 for snapshots E-5
 for virtual disks E-4

Lifeguard E-6

Microsoft Exchange E-8

preventing full pools E-6

POOL command A-8

pools

- adding storage units to 6-9, 10-5, A-12
 - benefits of 1-9
 - capacity of 2-4, 10-5, A-16
 - creating 6-2
 - creating in a cluster 6-5
 - defined Glossary-4
 - deleting 6-15, A-15
 - deleting in a cluster 6-1, 6-16
 - free space 8-13, E-3
 - getting full 6-5, 12-4, 12-5, E-6
 - maximum number of E-3
 - maximum number of snapshots E-5
 - maximum number of virtual disks 7-2, E-4
 - Microsoft Exchange pools D-2
 - monitoring free space 6-5, D-2
 - naming 6-4
 - properties 6-12
 - reconstructing 12-11
 - setting policies for E-3, E-6
 - showing information about 6-12, A-16
 - specifying free space E-3
 - storage unit requirements 2-2
 - supported disks 6-2
 - troubleshooting free space 12-9
 - unsupported disks 6-3
- possible owners
- defined Glossary-4
- preferred node
- defined Glossary-4
- privileges 4-12
- properties
- of network disks 9-17
 - of snapshots 8-12
 - of virtual disks 7-8

Q

quorum

- defined Glossary-4

quorum disk

- defined Glossary-4
- using a served network disk 9-15, C-25

R

RAID

- defined Glossary-4
- software RAID, not supported 2-3

reconstructing a pool 12-11

REGISTER command C-8

registered clients

- adding 9-5, 9-7, 9-8, C-9
- removing 9-5, 9-7, 9-8, C-11
- showing 9-8, C-12

Registry Editor E-2

removing

- computers from registered client lists 9-5, 9-7, C-11
- IP subnets from client configurations 9-13
- snapshots from pools A-26
- virtual disks from pools A-35

Restore from Snapshot wizard 11-2

restoring

- Restore from Snapshot wizard 11-2
- solutions 8-11
- virtual disks 7-3

restricting access to network disks 9-8

restrictions

- on network disk configurations 2-9
- on use of virtual disks and snapshots 2-3

Results pane of MMC 4-3

right-click menus 4-4

S

SCE Connected Disk resources 4-13, 12-10

- creating C-24
- deleting C-27

- SCE Pool resources 4-13, 12-7
 - creating A-10
 - deleting 6-16, A-15
 - in a pool 6-5
- SCE Served Disk resources 4-13
 - creating C-16
- SCE Storage Unit resources 4-13
 - deleting 6-16, A-15
 - in a pool 6-5, 6-11
- scheduled tasks
 - modifying 11-4
 - using Compaq Batch Scheduler 11-5
 - viewing 11-4, 11-5
- scheduling
 - automatic snapshots 11-3
- Scope pane of MMC 4-3
- SCSI
 - defined Glossary-5
- security
 - controlling which clients connect to your served network disks 9-8, C-8
 - privileges for snap-ins and commands 4-12
 - SWVR Admins group 4-12
- segments
 - and pool free space 12-9
 - copying out 12-9
 - defined Glossary-5
 - explained 8-2
 - specifying size 6-4
- SERVE command C-14
- server software 2-9
- serving network disks
 - disconnecting 9-21
 - disconnecting in a cluster 9-22
 - in a cluster 9-5, 9-7
 - other disks 9-5
 - restrictions 2-9
 - snapshots 9-4
 - starting 9-4, C-15
 - stopping 9-22, C-18
 - virtual disks 9-4
- shares
 - comparison with network disks 9-3
 - using with network disks 9-2
- shortening commands 4-11
- showing
 - available drive letters A-2, C-2
 - information about connected network disks 9-20
 - information about network disks 9-17, C-19, C-29
 - information about pools 6-12, A-16
 - information about snapshots A-28
 - information about virtual disks A-41
 - IP subnet configurations for clients 9-13
 - registered client lists 9-8, C-12
- size
 - see capacity
- SMARTSNAP utility B-2
- snap-ins
 - adding 4-3, 4-5
 - defined Glossary-5
 - explained 4-2
 - extensions 4-2
 - Network Disk Manager 4-2
 - privileges required 4-12
 - Snapshot Manager 4-2
 - using 4-1
- SnapMgr commands
 - abbreviating 4-11
 - how to use 4-8
 - managing remote computers and clusters 4-10
 - privileges required 4-12
 - summary of A-1
 - upper or lower case 4-11
 - DRIVES A-2
 - MANAGE A-4
 - POOL A-8
 - SNAPSHOT A-20
 - VIRTUALDISK A-32
- SNAPSHOT command A-20
- Snapshot for Backup wizard 8-10, 11-2
- Snapshot Manager snap-in 4-2

- Snapshot Planner 2-5
 - Snapshot Watchdog wizard 11-2
 - snapshots
 - and disk structure corrupt errors 12-7
 - benefits of 1-9
 - creating automatically 8-5
 - defined Glossary-5
 - deleting 8-13, A-26
 - deletion policies E-6, E-7
 - disk space used by 2-4, 2-5, 6-9, 7-7, 8-2, 8-13, A-26, A-29
 - displaying information 8-11
 - explained 8-2 to 8-3
 - keeping 8-5
 - mapping drive letters to 8-4, 8-5, A-23
 - maximum number in a pool E-5
 - modifying schedules 11-4
 - names of Glossary-5
 - naming 8-4
 - number of 8-3, E-5
 - of disks 8-4
 - planning 2-4, 2-5
 - policies for deleting E-7
 - properties 8-11, 11-4, 11-5
 - restrictions 2-2, 2-3
 - scheduling tasks 11-3, 11-5
 - serving as network disks 9-4
 - setting policies for E-5
 - showing information about A-28
 - unmapping drive letters from 8-6, A-25
 - using for backup 8-7, 8-10
 - using to restore virtual disks 7-3, 8-10
 - SNMP agent 1-5
 - SNMP Support 1-5
 - SQL 1-6
 - stalled I/Os to a network disk 12-2, C-25
 - standalone computers
 - defined Glossary-6
 - statistics
 - I/Os to network disks 9-19, C-20
 - I/Os to pools 6-14, A-17
 - stopping serving network disks 9-22, C-18
 - storage capacity
 - increasing 1-5
 - storage pools
 - see pools
 - storage units
 - adding in a cluster 6-11
 - adding to pools 2-2, 6-9, A-12
 - characteristics 2-2, 6-2
 - creating a pool from 6-2
 - defined Glossary-6
 - requirements for inclusion in a pool 6-2
 - showing which are in a pool 6-13, A-18
 - size limits 6-4
 - supported disks 6-2
 - stripe sets, host-based
 - not supported 2-3, 2-10
 - SWVR Admins group 4-12
 - system event log
 - failed to bring pool online 12-7
 - lost connection to network disk C-25
 - pool 70% full 6-6
 - pool 95% full 6-6
 - reconnected to network disk C-25
- T**
- Task Automation Wizards 11-2
 - task scheduling 11-1 to 11-6
 - methods 11-2
 - modifying 11-4, 11-5
 - TCP/IP 9-3
 - terms Glossary-1
 - testing applications 1-10
 - troubleshooting 12-1
 - tuning tips for network disks 9-16
- U**
- unmapping drive letters
 - from snapshots 8-6, A-25
 - from virtual disks 7-5, A-43
 - USE command C-22
 - user interfaces 1-7

V

View menu 4-4

virtual disks

- allocation unit size 7-3, E-4

- benefits of 1-9

- capacity 1-2, 1-3, 7-2

- creating 7-2

- defined Glossary-6

- defragmenting 7-7

- deleting 7-10, A-35

- disk space used by 2-4, A-42

- features 1-2

- growing 10-3, 10-5 to 10-8, A-38

- in clusters 7-2

- installing cluster-aware applications 1-2

- mapping drive letters to 7-1, 7-2, 7-4

- maximum number in a pool 7-2, E-4

- maximum number of snapshots 8-3, E-5

- names of Glossary-6

- naming 7-2

- partitioning and formatting 7-3, 7-6, 12-11, A-37

- planning 2-4

- policies E-4

- properties 7-8, 11-4, 11-5

- restoring 7-3

- restrictions 2-2, 2-3

- scheduling tasks 11-3, 11-5

- serving as network disks 9-4

- setting policies for E-4

- showing information about 7-7, A-41

- unmapping drive letters from 7-5, A-43

VIRTUALDISK command A-32

volume label 8-5

volume sets

- not supported 2-3, 2-10

volumes

- creating adjacent space 10-9

- data 10-20

- growing 1-5, 10-8, 10-15, 10-17

W

website, Compaq xi

Windows 2000

- installation requirement 3-1

Windows Explorer

- growing disks 10-22

Windows NT

- installation requirement 3-1

Windows Scheduled Tasks 11-2, 11-4

Workday Snapshot wizard 11-2

