

system
administrator's
guide

hp OpenView storage virtual replicator

Product Version: 4.0

Third Edition (April 2003)

Part Number: AA-RENGG-TE

This guide provides system administrators with information for installing, planning, and using Virtual Replicator for storage virtualization, online volume growth, and snapshots.



i n v e n t

© 2003 Hewlett-Packard Company

Hewlett-Packard Company makes no warranty of any kind with regard to this material, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. Hewlett-Packard shall not be liable for errors contained herein or for incidental or consequential damages in connection with the furnishing, performance, or use of this material.

This document contains proprietary information, which is protected by copyright. No part of this document may be photocopied, reproduced, or translated into another language without the prior written consent of Hewlett-Packard. The information contained in this document is subject to change without notice.

Compaq Computer Corporation is a wholly-owned subsidiary of Hewlett-Packard Company.

Adobe® and Acrobat® are trademarks of Adobe Systems Incorporated.

Microsoft®, MS-DOS®, MS Windows®, and Windows® and are U.S. registered trademarks of Microsoft Corporation.

Hewlett-Packard Company shall not be liable for technical or editorial errors or omissions contained herein. The information is provided "as is" without warranty of any kind and is subject to change without notice. The warranties for Hewlett-Packard Company products are set forth in the express limited warranty statements for such products. Nothing herein should be construed as constituting an additional warranty.

Printed in the U.S.A.

OpenView Storage Virtual Replicator System Administrator's Guide
Third Edition (April 2003)
Part Number: AA-RENGG-TE



contents

| | |
|--|-----------|
| About this Guide | 13 |
| Overview | 14 |
| Intended Audience | 14 |
| Related Documentation | 14 |
| Conventions | 15 |
| Document Conventions | 15 |
| Text Symbols | 16 |
| Equipment Symbols | 16 |
| Rack Stability | 18 |
| Getting Help | 19 |
| HP Technical Support | 19 |
| HP Storage Website | 19 |
| HP Authorized Reseller | 19 |
| 1 Introduction | 21 |
| Features | 22 |
| Storage Pooling | 22 |
| Virtual Disks | 22 |
| Snapshots | 23 |
| Online Volume Growth | 24 |
| Cluster Support | 25 |
| SNMP Support | 25 |
| Support for Windows 2000 Tools | 26 |
| Performance Monitor Support | 26 |
| Mounted Volume Support | 26 |
| Disk Defragmenter Support | 27 |
| Interoperability with Other Storage Management Tools | 27 |
| User Interface Options | 28 |
| Benefits | 29 |

| | | |
|----------|---|-----------|
| 2 | Installation | 33 |
| | Prerequisites | 34 |
| | Software Requirements | 34 |
| | Hardware Requirements | 34 |
| | License Requirements | 35 |
| | Components | 36 |
| | Installing Snapshot Planner | 37 |
| | Installing Virtual Replicator | 38 |
| | Setting Up Virtual Replicator Licensing | 40 |
| | Installing FLEXlm on a Dedicated License Server | 42 |
| | Installing in a Cluster | 43 |
| | Upgrading from Previous Versions | 44 |
| | Correcting Drive Letter Mapping After an Upgrade | 45 |
| | Upgrading to Virtual Replicator 4.0 on a Cluster | 45 |
| | Upgrading and Migrating to Windows Server 2003 | 46 |
| | Upgrading from an Evaluation Copy Without Rebooting | 47 |
| | Migrating Virtual Replicator to a Cluster | 48 |
| | Installing Adobe Acrobat Reader | 49 |
| 3 | Planning | 51 |
| | Pools, Virtual Disks, and Snapshots | 52 |
| | Storage Units in Pools | 53 |
| | Pool Capacity | 55 |
| | Rules for Using Virtual Disks and Snapshots | 56 |
| | Disk Space Requirements for Virtual Disks and Snapshots | 57 |
| | Using the Snapshot Planner | 58 |
| | Disk Space Requirements for Pools | 59 |
| | Advantages and Disadvantages of Large Pools | 60 |
| | Naming Conventions | 62 |
| | Existing Backup Tools | 63 |
| | Incremental Backup Support | 64 |
| 4 | Getting Started with Virtual Replicator | 65 |
| | Using the Microsoft Management Console Interface | 66 |
| | Understanding Snap-ins | 66 |
| | Snap-in Panes | 67 |
| | Performing Virtual Replicator Tasks | 67 |
| | Adding a Snap-in | 68 |
| | Managing Remote Computers | 72 |

| | |
|---|------------|
| Using the Command Line Interface | 73 |
| Using Commands to Manage Remote Computers | 73 |
| At the SnapMgr Command Prompt | 74 |
| At the Windows 2000/2003 Command Prompt | 74 |
| Abbreviating Commands | 74 |
| Using Uppercase or Lowercase in Commands | 75 |
| Using Virtual Replicator with Other Windows Tools | 76 |
| Administering Security and Privileges | 77 |
| Managing Other Domains | 77 |
| Using Cluster Resources | 78 |
| 5 Virtual Replicator - A Guided Tour | 79 |
| Creating a Pool | 80 |
| Creating a Virtual Disk in a Pool | 85 |
| Creating a Snapshot of a Virtual Disk | 91 |
| 6 Managing Pools | 97 |
| About Pools | 98 |
| Pools in a Cluster | 99 |
| Specifying Possible Owners for Pool Resources | 99 |
| Additional Storage Units in a Pool | 100 |
| About Importing Storage Units | 100 |
| About Removing Storage Units from Pools | 102 |
| Understanding Pool Free Space | 103 |
| Methods for Monitoring Pool Free Space | 103 |
| Using Replication Manager | 104 |
| Understanding Delspace | 105 |
| Virtual Disk Delspace | 105 |
| Snapshot Delspace | 106 |
| Delspace of Single and Multiple Snapshots | 106 |
| Information About Pools | 108 |
| About Pool Deletion | 112 |
| Pool Deletion in a Cluster | 112 |
| 7 Managing Virtual Disks | 113 |
| About Virtual Disks | 114 |
| Virtual Disks in a Cluster | 114 |
| About Drive Letter Mapping and Unmapping | 115 |
| Drive Letter Mapping in a Cluster | 115 |

| | |
|--|------------|
| Virtual Disk Formatting | 116 |
| Virtual Disk Defragmenting | 117 |
| Virtual Disk Properties | 118 |
| Using Virtual Replicator | 118 |
| Using My Computer. | 119 |
| Methods for Restoring Virtual Disks | 121 |
| The SnapBack Method of Restoring Virtual Disks | 121 |
| Example of Using SnapBack. | 122 |
| Other Methods for Restoring Virtual Disks | 123 |
| About Deleting Virtual Disks | 124 |
| About Growing Virtual Disks. | 125 |
| 8 Managing Snapshots | 127 |
| About Snapshots | 128 |
| Disk Space Usage by Snapshots | 128 |
| Best Practices for Using Snapshots | 130 |
| Snapshots in a Cluster | 130 |
| Automatic Snapshots | 131 |
| Snapshot Drive Letters | 132 |
| Snapshots for Backups | 133 |
| Snapshots for Restoring Virtual Disks. | 133 |
| Other Backup and Restore Solutions | 133 |
| Snapshot Properties | 134 |
| About Deleting Snapshots | 136 |
| 9 Online Volume Growth | 137 |
| Online Volume Growth for Basic and Virtual Disks | 138 |
| Understanding Basic, Virtual, and Dynamic Disks | 138 |
| Basic Disks | 138 |
| Virtual Disks | 139 |
| Dynamic Disks | 140 |
| Online Volume Growth Planning. | 140 |
| About Growing Virtual Disks. | 142 |
| About Growing Basic Disks | 143 |
| Basic Disk Preparation. | 143 |
| Adjacent Space on Basic Disks | 144 |
| Adjacent Space Through Deletion | 144 |
| Adjacent Space through Logical Drive Extension | 144 |
| Example of Online Volume Growth for Basic Disks. | 145 |

| | |
|--|------------|
| Configuration | 146 |
| Upgrading the Boot Volume | 147 |
| Using the ACU to Extend a Logical Drive | 147 |
| Using Virtual Replicator to Grow the Boot Volume | 148 |
| Online Volume Growth in Windows Explorer | 150 |
| 10 Scheduling Tasks | 151 |
| Task Scheduling Wizards Overview | 152 |
| 11 Troubleshooting | 153 |
| General Troubleshooting Questions | 154 |
| Incorrect Drive Letters After an Upgrade or Cluster Failover | 155 |
| Lost Delayed Write Errors | 156 |
| Applications Return Failed Write Errors | 157 |
| Accidentally Deleted a Virtual Disk | 158 |
| Pool Free Space Fell by More Than the Amount of Data Written | 159 |
| Reformatting a Virtual Disk Does Not Work | 160 |
| Reconstructing a Pool | 161 |
| Pools Not Coming Online | 162 |
| Prolonged Startup Times | 163 |
| Virtual Replicator Names Are Already in Use | 164 |
| Tracking Errors | 165 |
| Virtual Replicator Driver Events | 165 |
| Virtual Replicator Lifeguard Events | 166 |
| A Virtual Replicator Policy Management | 171 |
| About Policies | 172 |
| Policies in the Windows Registry | 173 |
| Pool Policies | 175 |
| Maximum Number of Pools (MaxPools) | 175 |
| Minimum Free Space for a Pool (MinFreespacePerPoolPercent) | 175 |
| Pool Segment Size (PoolSegmentSize) | 176 |
| Virtual Disk Policies | 177 |
| Maximum Number of Virtual Disks in a Pool (MaxVirtualPerPool) | 177 |
| Virtual Disk Allocation Unit Size (NTFSAllocationUnit) | 177 |
| Snapshot Policies | 179 |
| Maximum Number of Snapshots in a Pool (MaxSnapshotPerPool) | 179 |
| Maximum Number of Snapshots of a Virtual Disk (MaxSnapshotPerFamily) | 179 |
| Policies for the Virtual Replicator Lifeguard | 180 |

| | | |
|----------|--|------------|
| B | Virtual Replicator Lifeguard Service | 183 |
| | Monitoring Pool Free Space | 184 |
| | Monitoring Pool Free Space in a Microsoft Exchange Environment | 185 |
| | Checking Pools After Shutdown | 185 |
| C | SnapMgr Commands | 187 |
| | DRIVES | 189 |
| | Syntax | 189 |
| | Description | 189 |
| | Examples | 189 |
| | MANAGE | 190 |
| | Manage a Remote Computer | 191 |
| | Syntax | 191 |
| | Description | 191 |
| | Examples | 191 |
| | Manage the Local Computer | 192 |
| | Syntax | 192 |
| | Description | 192 |
| | Examples | 192 |
| | Show the Managed Computer | 193 |
| | Syntax | 193 |
| | Description | 193 |
| | Examples | 193 |
| | POOL | 194 |
| | Summary | 194 |
| | Create a Pool | 195 |
| | Syntax | 195 |
| | Description | 195 |
| | In a cluster | 196 |
| | Examples | 197 |
| | Add a Storage Unit to a Pool | 198 |
| | Syntax | 198 |
| | Description | 198 |
| | Examples | 199 |
| | Import a Storage Unit into a Pool | 200 |
| | Syntax | 200 |
| | Description | 200 |
| | Example | 201 |
| | Delete a Pool | 202 |

| | |
|--|-----|
| Syntax | 202 |
| Description. | 202 |
| Examples | 202 |
| Show Pools. | 203 |
| Syntax | 203 |
| Description. | 203 |
| Summary information | 203 |
| Full information | 204 |
| Showing I/O statistics | 204 |
| Showing the storage units | 205 |
| Examples | 205 |
| SNAPSHOT | 207 |
| Summary | 207 |
| Create a Snapshot. | 208 |
| Syntax | 208 |
| Description. | 208 |
| Examples | 209 |
| Map a Drive Letter to a Snapshot | 210 |
| Syntax | 210 |
| Description. | 210 |
| Examples | 210 |
| Unmap the Drive Letter from a Snapshot | 211 |
| Syntax | 211 |
| Description. | 211 |
| Examples | 211 |
| Delete a Snapshot. | 212 |
| Syntax | 212 |
| Description. | 212 |
| Examples | 212 |
| Show Snapshots | 213 |
| Syntax | 213 |
| Description. | 213 |
| Summary information | 213 |
| Full information | 214 |
| Examples | 214 |
| Assign a Mount Point to a Snapshot | 215 |
| Syntax | 215 |
| Description. | 215 |

| | |
|---|-----|
| Example | 215 |
| Remove a Mount Point from a Snapshot | 216 |
| Syntax | 216 |
| Description. | 216 |
| Example | 216 |
| List the Mount Points of a Snapshot | 217 |
| Syntax | 217 |
| Description. | 217 |
| Example | 217 |
| UNITS. | 218 |
| Syntax | 218 |
| Description. | 218 |
| Example | 219 |
| UTILITY. | 220 |
| Syntax | 220 |
| Description. | 220 |
| Examples | 221 |
| VIRTUALDISK | 222 |
| Summary | 222 |
| Assign a Mount Point to a Virtual Disk. | 223 |
| Syntax | 223 |
| Description. | 223 |
| Example | 223 |
| Create a Virtual Disk | 224 |
| Syntax | 224 |
| Description. | 224 |
| Examples | 224 |
| Delete a Virtual Disk | 225 |
| Syntax | 225 |
| Description. | 225 |
| Examples | 225 |
| Format a Virtual Disk. | 226 |
| Syntax | 226 |
| Description. | 226 |
| Examples | 226 |
| Grow a Virtual Disk. | 227 |
| Syntax | 227 |
| Description. | 227 |

| | |
|--|------------|
| Example | 227 |
| List Mounts Points of a Virtual Disk | 228 |
| Syntax | 228 |
| Description. | 228 |
| Example | 228 |
| Map a Drive Letter to a Virtual Disk | 229 |
| Syntax | 229 |
| Description. | 229 |
| Examples | 229 |
| Remove a Mount Point from a Virtual Disk | 230 |
| Syntax | 230 |
| Description. | 230 |
| Example | 230 |
| Restore a Virtual Disk | 231 |
| Syntax | 231 |
| Description. | 231 |
| Examples | 231 |
| Show Virtual Disks | 232 |
| Syntax | 232 |
| Description. | 232 |
| Summary information | 232 |
| Full information | 233 |
| Examples | 233 |
| Unmap the Drive Letter from a Virtual Disk | 234 |
| Syntax | 234 |
| Description. | 234 |
| Examples | 234 |
| D SmartSnap Command | 235 |
| SMARTSNAP | 236 |
| Syntax. | 236 |
| Description | 236 |
| Example | 237 |
| Index | 239 |



about this guide

This system administrator guide provides information to help you:

- Plan your Virtual Replicator implementation.
- Understand all of Virtual Replicator's storage management features.
- Get started using the Virtual Replicator software.

About this Guide topics include:

- [Overview](#), page 14
- [Conventions](#), page 15
- [Rack Stability](#), page 18
- [Getting Help](#), page 19

Overview

This section covers the following topics:

- [Intended Audience](#)
- [Related Documentation](#)

Intended Audience

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

- Management of Microsoft Windows 2000 and 2003 servers and workstations
- Configuration and management of Microsoft Cluster Server (MSCS) clusters

Related Documentation

In addition to this guide, HP provides corresponding information:

- Online help
- Installation reference card
- Commands reference card
- Planning chart

Conventions

Conventions consist of the following:

- [Document Conventions](#)
- [Text Symbols](#)
- [Equipment Symbols](#)

Document Conventions

The document conventions included in [Table 1](#) apply in most cases.

Table 1: Document Conventions

| Element | Convention |
|---|---|
| Cross-reference links | Blue text: Figure 1 |
| Key and field names, menu items, buttons, and dialog box titles | Bold |
| File names, application names, and text emphasis | <i>Italics</i> |
| User input, command and directory names, and system responses (output and messages) | Monospace font COMMAND NAMES are uppercase monospace font unless they are case sensitive |
| Variables | <monospace, italic font> |
| Website addresses | Blue, underlined sans serif font text: http://www.hp.com |

Text Symbols

The following symbols may be found in the text of this guide. They have the following meanings.



WARNING: Text set off in this manner indicates that failure to follow directions in the warning could result in bodily harm or death.



Caution: Text set off in this manner indicates that failure to follow directions could result in damage to equipment or data.

Note: Text set off in this manner presents commentary, sidelights, or interesting points of information.

Equipment Symbols

The following equipment symbols may be found on hardware for which this guide pertains. They have the following meanings.



Any enclosed surface or area of the equipment marked with these symbols indicates the presence of electrical shock hazards. Enclosed area contains no operator serviceable parts.

WARNING: To reduce the risk of personal safety from electrical shock hazards, do not open this enclosure.



Any RJ-45 receptacle marked with these symbols indicates a network interface connection.

WARNING: To reduce the risk of electrical shock, fire, or damage to the equipment, do not plug telephone or telecommunications connectors into this receptacle.



Any surface or area of the equipment marked with these symbols indicates the presence of a hot surface or hot component. Contact with this surface could result in injury.

WARNING: To reduce the risk of personal safety from a hot component, allow the surface to cool before touching.



Power supplies or systems marked with these symbols indicate the presence of multiple sources of power.

WARNING: To reduce the risk of personal safety from electrical shock, remove all power cords to completely disconnect power from the power supplies and systems.



Any product or assembly marked with these symbols indicates that the component exceeds the recommended weight for one individual to handle safely.

WARNING: To reduce the risk of personal safety or damage to the equipment, observe local occupational health and safety requirements and guidelines for manually handling material.

Rack Stability

Rack stability protects personnel and equipment.



WARNING: To reduce the risk of personal safety or damage to the equipment, be sure that:

- The leveling jacks are extended to the floor.
 - The full weight of the rack rests on the leveling jacks.
 - In single rack installations, the stabilizing feet are attached to the rack.
 - In multiple rack installations, the racks are coupled.
 - Only one rack component is extended at any time. A rack may become unstable if more than one rack component is extended for any reason.
-

Getting Help

If you still have a question after reading this guide, contact an HP authorized service provider or access our website: <http://www.hp.com>.

HP Technical Support

Telephone numbers for worldwide technical support are listed on the following HP website: <http://www.hp.com/support/>. From this website, select the country of origin.

Note: For continuous quality improvement, calls may be recorded or monitored.

Be sure to have the following information available before calling:

- Technical support registration number (if applicable)
- Product serial numbers
- Product model names and numbers
- Applicable error messages
- Operating system type and revision level
- Detailed, specific questions

HP Storage Website

The HP website has the latest information on this product, as well as the latest drivers. Access storage at: <http://www.hp.com/country/us/eng/prodserv/storage.html>. From this website, select the appropriate product or solution.

HP Authorized Reseller

For the name of your nearest HP authorized reseller:

- In the United States, call 1-800-345-1518
- In Canada, call 1-800-263-5868
- Elsewhere, see the HP website for locations and telephone numbers: <http://www.hp.com>.

Introduction



1

HP OpenView Storage Virtual Replicator provides system administrators with advanced, centralized storage virtualization and management capabilities in Microsoft Windows 2000 and Windows Server 2003 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 MSCS clusters.

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

- [Features](#), page 22
 - [Storage Pooling](#), page 22
 - [Virtual Disks](#), page 22
 - [Snapshots](#), page 23
 - [Online Volume Growth](#), page 24
 - [Cluster Support](#), page 25
 - [SNMP Support](#), page 25
 - [Support for Windows 2000 Tools](#), page 26
 - [Interoperability with Other Storage Management Tools](#), page 27
 - [User Interface Options](#), page 28
- [Benefits](#), page 29

Features

Virtual Replicator allows you to:

- Group hardware arrays or physical disks to form a large **pool** of storage.
- Divide the pool into **virtual disks** of any size, up to 2 TB (terabytes).
- Make instant copies, called **snapshots**, of the virtual disks.
- Use the virtual disks and snapshots on the **local** computer.

You can access Virtual Replicator tools through either the Microsoft Management Console (MMC) or a command line interface (CLI).

These and other Virtual Replicator features, such as support for cluster configurations, SNMP, and Microsoft Windows Performance Monitor, are described in the sections that follow.

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 2000/Windows Server 2003 has direct access can be used in a pool. In addition to standard single-spindle disks, you can use controller-based, fault-tolerant disk arrays, such as HP 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 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. 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 match the sizes of the virtual disks with 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 1-TB database,

you can combine several disks or HP StorageWorks RAID arrays in a single pool, and create a 1-TB virtual disk that spans that physical storage. The size of the virtual disks you create can range from 10 MB to 2 TB, depending on the pool's free space, its segment size, and the policies in effect for the pool.

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. You can read from snapshots and write to them, and you do not have to interrupt your users to create them. You can create snapshots while users are reading from and writing to the original virtual disk.

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. Making a snapshot is like taking a picture of every byte of data on the original virtual disk at a single instant.

Initially, a snapshot uses almost no disk space from its pool. The snapshot starts to use disk space only when you write data to either the original virtual disk or the snapshot itself.

[Figure 1](#) shows a storage pool that is built from two 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.

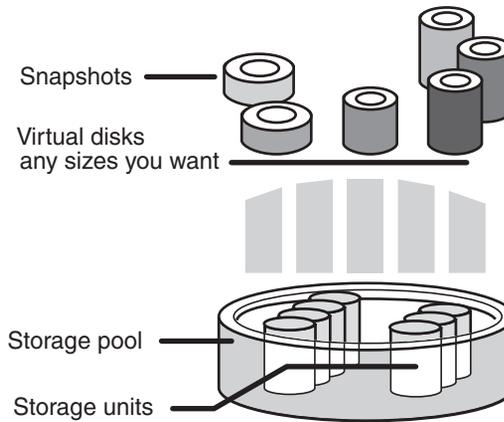


Figure 1: A Storage Pool with Virtual Disks and Snapshots

Snapshots are useful 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.

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

See Chapter 8, [Managing Snapshots](#), for more information.

Virtual Replicator also provides wizards to automatically schedule unattended creation, deletion, and backup of snapshots. See Chapter 10, [Scheduling Tasks](#), for information on these wizards.

Online Volume Growth

Virtual Replicator lets you 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 function allows you to grow a volume 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 9, [Online Volume Growth](#), 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 78 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\HP\OPENVIEW 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.

Support for Windows 2000 Tools

Virtual Replicator offers integration with the following Windows 2000 tools and features:

- Performance (Performance Monitor)
- Mounted volumes
- Disk Defragmenter

Performance Monitor Support

Virtual Replicator supports the Windows Performance tool. This tool shows the behavior of virtual disks and snapshots using disk performance counters, such as Disk Read Bytes/sec and Disk Write Bytes/sec. Performance counters for virtual disks and snapshots are listed under the **PhysicalDisk** or **LogicalDisk** performance object.

In addition, Virtual Replicator adds its own objects to the list of objects that can be monitored. These objects are named **SWVR_Pool**, **SWVR_VirtualDisk**, and **SWVR_Snapshot**. Using the special counters listed under these objects, you can monitor performance activity specific to Virtual Replicator, for example, available free space in a pool or the Delspace of a snapshot.

Mounted Volume Support

Virtual disks created with Virtual Replicator can be used as Windows 2000 volume mount points. The Windows 2000 Disk Management tool allows you to set up directories on a volume as mount points for other volumes. Users and applications can refer to a mounted volume by its mount path, rather than by a drive letter. Mount points give you access to more drives because you are not restricted to the 26-drive letter limit.

Note: Volume mount points are not supported on cluster systems.

The mount point feature can be effectively used for managing virtual disks and snapshots. For example, you can assign a virtual disk as a mount point and mount the snapshots to the virtual disk. The snapshots do not use up drive letters and are easy to manage.

Note that several Virtual Replicator features require virtual disks to have a mapped drive letter. Incremental backup support, online volume growth, and virtual disk formatting all require drive letters. Therefore, HP does not recommend mounting virtual disks to other volumes.

For more information about mounted volumes, see the Microsoft Windows 2000 documentation.

Disk Defragmenter Support

Virtual Replicator is fully integrated with the Windows Disk Defragmenter tool. You can defragment virtual disks to increase I/O performance in much the same way as defragmenting a physical disk. The Disk Defragmenter operation finds and consolidates fragmented files and folders on the volume. The tool also consolidates your free space, making it less likely that new files will be fragmented. As a result, your system can gain access to your files and folders and save new ones more efficiently. See [Virtual Disk Defragmenting](#) on page 117 for more information.

Interoperability with Other Storage Management Tools

Virtual Replicator operates jointly with other storage applications from HP, such as HP StorageWorks Secure Path and HP StorageWorks Command Console (SWCC), as well as storage applications from 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 backup tools, such as Computer Associates ArcServe and VERITAS Backup Exec, as well as database and messaging applications, such as Microsoft SQL Server and Microsoft Exchange. To read white papers that discuss these applications, refer to the website <http://h18000.www1.hp.com/products/sanworks/vr/index.html>.

User Interface Options

You can use Virtual Replicator in the following ways:

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

Each of these methods provides you with full capabilities to monitor and manage your storage resources. See Chapter 4, [Getting Started with Virtual Replicator](#), for more information on the interface options.

Benefits

Here are some ways that typical enterprises can benefit from Virtual Replicator.

- Optimize disk storage by using storage pooling/virtualization.

You can optimize disk storage by centralizing multiple logical units and hardware HP 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.

- Add storage capacity and grow volumes without rebooting.

Virtual Replicator allows you 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 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 make it available to your application servers in the following ways:

- 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 in Chapter 8, [Managing Snapshots](#).

- 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. Using snapshots 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.

For 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, create a snapshot of the virtual disk that holds the database data, so that if the batch job fails partway 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. 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.

Installation

2

This chapter includes the following topics:

- [Prerequisites](#), page 34
- [Components](#), page 36
- [Installing Snapshot Planner](#), page 37
- [Installing Virtual Replicator](#), page 38
- [Setting Up Virtual Replicator Licensing](#), page 40
- [Installing FLEXlm on a Dedicated License Server](#), page 42
- [Installing in a Cluster](#), page 43
- [Upgrading from Previous Versions](#), page 44
- [Migrating Virtual Replicator to a Cluster](#), page 48
- [Installing Adobe Acrobat Reader](#), page 49

Prerequisites

Before you install Virtual Replicator V4.0, make sure your system meets all of the specified requirements.

Software Requirements

The following software is required:

- **Either** of the following operating systems:
 - Windows 2000 – Professional, Server, or Advanced Server, Service Pack 2 or 3
 - Windows Server 2003

Note: Of the versions listed for Windows 2000, only Advanced Server supports clustering.

- Microsoft Internet Explorer, Version 5.01 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/>.
- Microsoft Management Console 1.1 or greater
MMC is included with the Windows 2000 and Windows Server 2003 operating systems.

Hardware Requirements

The hardware prerequisites conform to Microsoft minimum requirements for the particular operating system, including:

- An Intel-architecture, 32-bit processor
- Minimum of 128 MB of memory; 256 MB recommended
- 100 MB of disk space for a full installation

License Requirements

Before you can use the Virtual Replicator software on one or more systems, you need to have a valid license. To acquire the license, follow the instructions on the License Key Retrieval Instruction Sheet included with the Virtual Replicator license kit (ordered separately).

For License Manager to accept a username and password, the Java Runtime Environment plugins (j2re-1_3_1_02-win.exe and plugin-1_1_3_005-win.exe) are required. The plugins can be obtained from <http://java.sun.com/products/>.

Note: The Virtual Replicator V4.0 license agreement and access key that are included in your license kit apply to the Virtual Replicator V4.0 software.

As explained in the instruction sheet, you obtain a license key by going to the licensing website and providing the Authorization ID printed on the sheet. Within 48 hours, you will receive an e-mail message from HP that contains the information you need to properly set up licensing and operate Virtual Replicator.

HP recommends that you begin the process of retrieving your license before installing Virtual Replicator. After you receive your license key, use the instructions in [Setting Up Virtual Replicator Licensing](#) on page 40 to complete the procedure.

Components

Version 4.0 of Virtual Replicator includes two installable components:

- Snapshot Planner

Use this tool to help predict how much disk space you need for your snapshots. Snapshot Planner is useful only *before* you create Virtual Replicator resources.

- Virtual Replicator

Use Virtual Replicator to create pools of storage, divide the pools into virtual disks, create snapshots of the virtual disks, and grow volumes dynamically.

Installing Snapshot Planner

If you want to use the optional Snapshot Planner tool, install and use it **before** you set up your pools and virtual disks.

To install Snapshot Planner:

1. Log on to a user account that has administrator privileges.
2. Insert the HP OpenView Storage Virtual Replicator installation CD into your CD-ROM drive. Installation begins automatically.

Note: Alternatively, you can double click SETUP.EXE from the root directory on the CD.

3. Click **Snapshot Planner** on the first screen. Follow the instructions provided by the setup wizard until you are finished.

HP recommends that you use and uninstall Snapshot Planner before installing Virtual Replicator. See [Using the Snapshot Planner](#) on page 58 for more information.

Installing Virtual Replicator

For new installations of Virtual Replicator, perform the following steps:

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

Note: Alternatively, you can double click SETUP.EXE from the root directory on the CD.

3. Click **Virtual Replicator** on the first screen.
The **License Agreement** screen is displayed.
4. Click the button to accept the terms of the license agreement, then click **Next**.
The **License Server Location** screen is displayed. Virtual Replicator uses the Globetrotter FLEXlm license server for management of VR licenses.
5. Choose whether to install the license server on the local computer or to connect to another computer that already has the license server installed. To install FLEXlm on a dedicated server, see [Installing FLEXlm on a Dedicated License Server](#) on page 42.
 - If you are a new user of the FLEXlm license server, click the button to install the license server on the local computer.
 - If the FLEXlm license server has been previously installed elsewhere, click the button to connect to another computer. Then, type the IP address or hostname of the other computer.
6. Click **Next**.
The **Setup Type** screen is displayed.
7. 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.
8. Continue through the wizard screens until you have finished.

Note: Restart your computer after completing the installation if you are instructed to do so.

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

After installing Virtual Replicator, you cannot use the product's features until you set up the required licenses. See the next section, [Setting Up Virtual Replicator Licensing](#), for more information.

Setting Up Virtual Replicator Licensing

Set up and manage Virtual Replicator licensing by using Storage License Manager (SLM). This tool is installed during the installation of Virtual Replicator. SLM controls the number of Virtual Replicator licenses, which ensures that all systems are licensed and helps users comply with the license agreement.

Storage License Manager allows software licenses to be available (or float) on the network instead of being tied to a particular machine. Users can more efficiently use fewer licenses by sharing them on the network. As the Administrator, you can control who uses the licenses and on what machine (or node) the licenses are available.

As explained in [License Requirements](#) on page 35, you receive the license key that allows you to set up licensing and operate Virtual Replicator in an e-mail message from HP.

After you receive the license key from HP, use SLM to create a license file. The license file allows SLM to implement the license properly. To start SLM:

1. Click **Start > Programs > SANworks Storage License Manager > Storage License Manager**.

The **Compaq Web-Based Management Account Login** window opens.

If you installed the license server on the local computer, continue with step 2.

If the license server is installed on another computer, connect to it by replacing **localhost** in the web browser **Address** box with the hostname or IP address of the other computer. Then continue with step 2.

2. Type **administrator** in both the **Name** and **Password** boxes, and click **OK**.
administrator is the default SLM username and password for new users. You can change them later.

The SLM web browser interface opens.

3. Click **Licenses** on the License Manager menu bar.

The **License Files** window opens.

4. Click **Add** to display the **Add License File** page.
5. Type a unique name for the license file, then click **Create**.
6. Click **Close** to return to the **License Files** page.
7. Select the name of the license file you just created from the list of current licenses and click **Modify**.

The **Contents of License File** window opens at the bottom of the page.

8. Open the e-mail message from HP that contains your license key text.
9. Cut and paste the license key text from the e-mail message into the **Contents of License File** window.
10. Click **Save**.
11. Click **Exit** on the menu bar.

Your system is now properly registered to operate Virtual Replicator.

You can use SLM to add, change, or delete license files at a later time. For more information about SLM tasks and features, see the Storage License Manager online help.

Installing FLEXIm on a Dedicated License Server

You can install the FLEXIm License Manager on a dedicated license server without having to install Virtual Replicator.

Note: You should install FLEXIm first before installing Virtual Replicator on the other servers.

Follow these steps:

1. Open the Virtual Replicator kit.
2. Locate and open the Storage Software folder.
3. Locate and open the Storage License folder.
4. Double-click the setup icon to install the FLEXIm License Manager. The install wizard opens.
5. Follow the wizard instructions. When prompted, select **Use this computer as the license server**.
6. When the installation is done, follow the steps in [Setting Up Virtual Replicator Licensing](#) on page 40.

Note: Enter the access key redeemed from the website using the authorization ID from the license kit.

7. Follow the steps in [Installing Virtual Replicator](#) on page 38. In step 5 of that procedure, follow the instructions in the second bulleted item.

Installing in a Cluster

The Virtual Replicator components are fully cluster aware. On a cluster, install Virtual Replicator on each node following the instructions in [Installing Virtual Replicator](#). Be sure to install exactly the same Virtual Replicator options on all nodes of a cluster. For example, if you choose the Custom Setup option and install the Online Volume Growth feature on one node, make sure you install it on all the nodes in the cluster.

You must set up licensing for both nodes in a cluster by:

- Installing Storage License Manager on both nodes or pointing both nodes to a remote SLM server.
- Obtaining a separate Virtual Replicator license for each node.

See [Setting Up Virtual Replicator Licensing](#) on page 40 for more information.

Upgrading from Previous Versions

If you are a current user of Virtual Replicator Version 2.5 and above, you can use the installation procedure described in [Installing Virtual Replicator](#). Before doing so, HP recommends you take note of the current drive letter mappings for your virtual disks and snapshots. This step is advised so that you can correct drive letter mappings after the upgrade. See [Correcting Drive Letter Mapping After an Upgrade](#) below for more information.

Note: If you want to upgrade from Virtual Replicator version 2.0 or below to Version 4.0, you must first uninstall Virtual Replicator manually and then use the procedure described in [Installing Virtual Replicator](#).

When upgrading, HP recommends that you **do not manually uninstall** Virtual Replicator Version 2.5 and above. As part of the upgrade, Version 4.0 preserves your storage environment and then uninstalls the earlier version of Virtual Replicator. You can continue using your current pools, virtual disks, and snapshots.

Note: Virtual Replicator Version 4.0 does not support network disk serving.

If you are serving virtual disks as network disks, do the following before upgrading to Version 4.0:

1. Back up the data on all network disks.
2. Disconnect all client computers using network disks.
3. Stop serving all network disks.
4. Install Virtual Replicator 4.0.

Correcting Drive Letter Mapping After an Upgrade

When upgrading from previous versions of Virtual Replicator on standalone or cluster systems, you may find that previously mapped virtual disks and snapshots no longer have the same drive letters.

To preserve drive letter mapping, HP recommends you take the following steps:

1. Before upgrading, take note of the current drive letter configuration for your virtual disks and snapshots.
2. After you upgrade, and after all virtual disks and snapshots have been remapped, make sure that the previously used drive letters are still available, that is, they are not assigned to non-VR objects.
3. Use the Restore Drives utility to restore the original mappings.

Depending on where Virtual Replicator is installed on your system, this command looks like the following:

```
C:\>PROGRAM FILES\HP\OPENVIEW VIRTUAL REPLICATOR 4.0\  
RESTOREDIVES.EXE
```

4. After running the utility, refresh the view in the Replication Manager MMC window (**Action > Refresh**).

All drive letters should be correctly mapped.

See [Incorrect Drive Letters After an Upgrade or Cluster Failover](#) on page 155 for more information about the Restore Drives utility.

Upgrading to Virtual Replicator 4.0 on a Cluster

Use the following procedure when upgrading from Virtual Replicator Version 2.5 and above to Version 4.0 on a cluster system.

In a two-node cluster with nodes A and B:

1. Move all cluster groups to node B and then install Virtual Replicator 4.0 on node A using the instructions in [Installing Virtual Replicator](#) on page 38.
2. Reboot node A and then move all of the cluster groups over to node A.
3. Using Cluster Administrator, confirm that any pre-existing pools are online on node A.
4. Install Virtual Replicator 4.0 on node B and reboot.
5. Using Cluster Administrator, confirm that any pre-existing pools are online on node B.

Upgrading and Migrating to Windows Server 2003

If you want to upgrade from Version 2.5 and above to Version 4.0 *and* migrate from Windows 2000 to Windows Server 2003 on a standalone computer, use the following procedure:

Note: This procedure is not supported on clusters. Also, any existing mount points will be lost during the migration from Windows 2000 to Windows Server 2003.

1. Back up your data.
2. Install Virtual Replicator 4.0 using the procedure described in [Installing Virtual Replicator](#) on page 38.
Version 4.0 preserves your storage environment and then uninstalls the previous version of Virtual Replicator.
3. Manually uninstall Virtual Replicator 4.0 from the Windows 2000 system.
This step is required to maintain your Virtual Replicator environment during the operating system migration.

4. Upgrade the system to Windows Server 2003.
5. Install Virtual Replicator 4.0 again.
6. Run the Namespace Recovery Utility, using the following command:

```
SNAPMGR.EXE UTILITY /RECOVER:NAMESPACE
```

This command updates the Windows Registry based on your earlier configuration. As a result, your Virtual Replicator environment is reconstructed. If you had custom policies specified in your previous configuration, these are **not** restored. Instead, the pre-set default policies are in place.

See [Virtual Replicator Names Are Already in Use](#) on page 164 for more information about the Namespace Recovery Utility.

7. Use the Restore Drives utility to restore the original drive letter mappings. Depending on where Virtual Replicator is installed on your system, this command looks like the following:

```
C:\>PROGRAM FILES\HP\OPENVIEW VIRTUAL REPLICATOR 4.0\  
RESTOREDIVES.EXE
```

See [Incorrect Drive Letters After an Upgrade or Cluster Failover](#) on page 155 for more information about the Restore Drives utility.

Upgrading from an Evaluation Copy Without Rebooting

If you have an evaluation copy of Virtual Replicator, you can upgrade it to a full retail licensed copy as follows:

1. Add the full use license kit to Storage License Manager.
2. Stop and start the Virtual Replicator Management service from either the command line or the desktop:

From the command line:

- Enter the following commands:

```
net stop esmgrs  
net start esmgrs
```

The new license is activated.

From the desktop:

- a. Right-click **My Computer** and select **Manage > Services & Applications > Services > Virtual Replicator Management Service**.
The **Properties** dialog box for the Management service is displayed.
- b. Click **Stop** and then click **Apply**.
- c. Click **Start** and then click **Apply**.
- d. Click **OK** to close the **Properties** dialog box.
The new license is activated.

Migrating Virtual Replicator to a Cluster

If you have installed and set up Virtual Replicator on a standalone system and later want to migrate to a cluster system, you need to uninstall and then reinstall Virtual Replicator in the cluster environment.



Caution: Before using the migration procedure, make sure you can successfully back up and restore the data in your Virtual Replicator environment.

Use the following procedure to configure Virtual Replicator so that it operates and fails over properly in a cluster:

1. Back up your data.
2. Record the names of pools, virtual disks, and snapshots on the standalone system.
3. Delete all pools, virtual disks, and snapshots.
4. Uninstall Virtual Replicator from the standalone system.
5. Install the MSCS software on systems A and B.
6. Reinstall Virtual Replicator on system A.
7. Install Virtual Replicator in the same way on system B.

Note: Each system must have a separate license to operate Virtual Replicator.

8. Recreate the same pools, virtual disks, and snapshots that you had on the standalone system, using the previously recorded names.
9. Restore data from your backup media to the pools, virtual disks, and snapshots.

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>.

Planning

3

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 52
- [Storage Units in Pools](#), page 53
- [Pool Capacity](#), page 55
- [Rules for Using Virtual Disks and Snapshots](#), page 56
- [Disk Space Requirements for Virtual Disks and Snapshots](#), page 57
- [Using the Snapshot Planner](#), page 58
- [Disk Space Requirements for Pools](#), page 59
- [Advantages and Disadvantages of Large Pools](#), page 60
- [Naming Conventions](#), page 62
- [Existing Backup Tools](#), page 63
- [Incremental Backup Support](#), page 64

Note: To help plan the Virtual Replicator configurations, use the planning charts in the Documentation folder on the Virtual Replicator distribution CD-ROM. The CD-ROM includes a blank chart to fill in ([Planning chart.pdf](#)), and an example of a completed chart ([Example planning chart.pdf](#)).

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 limits are allowed for virtual disks and snapshots:

- Up to 8 virtual disks in each pool.

The capacity of each virtual disk can be between 10 MB and 2 TB. The capacity is limited by the pool's free space, its segment size, and by the policies in effect for the pool.

- Up to 12 snapshots in each family.

A family consists of a virtual disk, its snapshots, and the snapshots of those snapshots. 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 8 storage units. They can have different capacities and be from different manufacturers.
- 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.
- Do not use disks that are part of volume sets, mirrorsets, or stripesets created using Disk Administrator or Disk Management. Use only controller-based units.

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

HP recommends that all the storage units in a pool have the same redundancy, read-write, and failure characteristics. For example, they should all be HP StorageWorks RAID 5 storagesets, or they should all be mirrorsets, 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 mirrorsets.

Note: Do not use Disk Administrator or Disk Management to create a mirrorset on a virtual disk (see [Rules for Using Virtual Disks and Snapshots](#) on page 56).

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. That is, if one of the pool's storage units becomes inaccessible due to a hardware fault, you lose the whole pool and all of its virtual disks and snapshots.

For example, you should not use a standard single-spindle disk and a mirrorset (RAID 1) in the same pool. Although the 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.

If you lose a pool, you have to reconstruct it. To do this, you 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.

Pool Capacity

A pool can contain up to 8 logical storage units, and each storage unit can have a capacity of up to 1 TB. Your system's hardware and software configurations determine the maximum size of the pools you create. For example, using the HP HSG80 controller and the ACS Version 8.6 firmware, you can configure storage units of 1 TB. The maximum pool you can create in this setting is 8 TB. Other configurations will yield different results.

The actual capacity of the 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 percent of the combined capacities of the storage units.

For example, if you create a pool from five 20-GB RAID storage sets, assume for planning purposes that the capacity of the pool is 90 GB. You can increase the capacity of a pool at any time by adding storage units to it. You can do this dynamically, while users are reading from and writing to the virtual disks and snapshots in the pool.

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 Management for Windows 2000 or Windows Server 2003 to create a volume set, mirror set, or stripe set on the virtual disk.

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

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, the snapshot does not use any space. The snapshot consumes space only when users modify the data stored on either the snapshot itself or its parent. The more data users 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 indicated by the incremental save set, 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 an optional tool called Snapshot Planner to help you predict how much space your snapshots will use. For best use of this tool:

1. Install Snapshot Planner separately, *before* installing Virtual Replicator.
To install, click **Snapshot Planner** on the installation startup screen.
2. Use Snapshot Planner to profile your current storage environment.
To run Snapshot Planner, click **Start > Programs > hp OpenView Storage Virtual Replicator > Snapshot Planner**.
3. Uninstall Snapshot Planner before you install Virtual Replicator. To uninstall:
 - a. Open the Windows Control Panel and double click **Add/Remove Programs**.
 - b. Click **OpenView Virtual Replicator Snapshot Planner** and click **Remove**.

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.

Snapshot Planner tracks actual reads and writes to your existing drives. It calculates 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.

The Snapshot Planner tool is useful if you plan to migrate all the data on an existing hard disk drive to a virtual disk of the same capacity. The tool 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, you might need to increase or decrease the values recommended by Snapshot Planner.

For more information, see the Snapshot Planner online help.

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, the applications will report failed writes immediately. You will realize right away that the pool is full.

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, do not try to free up space by deleting files. Deleting files causes segments of data to be copied out to snapshots, which uses up even more free space in the pool and can cause additional data loss.

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 the pool could run out of disk space and become full.

In addition, keeping snapshots for long periods of time can contribute to pools becoming full, and can also increase the time needed to restart the system. Keep snapshots only as long as needed and delete unnecessary snapshots, whenever possible.

When snapshots are present, monitor the pool's free space regularly and make sure that it does not become full. For more information, see [Methods for Monitoring Pool Free Space](#) on page 103.

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 share 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 HP 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 striping to balance loads across the different storage units in a pool. However, when a pool has multiple units, Virtual Replicator uses the next available LUN when creating each new virtual disk.
- 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 configuration saves 50 percent 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 percent 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.

Naming Conventions

Every pool, virtual disk, and snapshot must have a unique 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. In addition, because Virtual Replicator automatically creates cluster groups and resources using the names you choose, a pool cannot have the same name as any existing cluster group or cluster resource. For example, if there is a cluster resource called Res1, do not create a pool named Res1.

To conform to Windows file naming conventions, avoid using the following characters in Virtual Replicator names:

\ / : * ? " < > | =

Existing Backup Tools

You can use your existing backup tools to back up virtual disks and snapshots. If your backup utility performs incremental backups, Virtual Replicator provides an additional support feature: Incremental Backup Support.

Incremental Backup Support operates by turning off the archive file attributes, or archive bits, of files on a virtual disk after one of its snapshots is backed up.

Incremental Backup Support

Some backup utilities use an archive bit to indicate whether or not a file has been backed up. When a file is backed up, the utility clears or turns “off” the file’s archive bit, indicating that the file does not need to be backed up again during the next incremental backup. If the file changes prior to the next backup, the bit will be turned “on” again, and the utility will back up the file during the next incremental backup.

When backing up a snapshot, the backup utility turns off the archive bit of each file on the snapshot. However, the utility does not clear the archive bits of the corresponding files on the snapshot's parent disk; therefore, they are still on. If a new snapshot is created from the same parent disk, the archive bit of every file on the new snapshot will also be on, even though the files might not have changed between snapshots.

The **on** status indicates to the backup utility that *all* of the new snapshot’s files need to be backed up in the next incremental operation, possibly causing duplicate backups.

The Incremental Backup Support feature allows normal incremental backups by turning off the archive bits of the files on the virtual disk. The next snapshots that are taken will then have the archive bits correctly set for incremental backups.

To enable Incremental Backup Support, you can specify one of the following:

- A date and time that the Incremental Backup Support operation will use as a reference point, along with a safety margin.

The archive bit will be turned off for files on the parent disk that are older than the specified date and time. The safety margin helps prevent the incorrect turning off of archive bits for files that have not been backed up.

- The name of a snapshot that was backed up, if the snapshot still exists. The archive bit will be turned off for files on the parent disk that correspond to the snapshot’s files.

You also have the option of creating a log file of Incremental Backup Support operations. Full instructions for using the Incremental Backup Support feature are available in the online help.

Getting Started with Virtual Replicator

4

This chapter provides an introduction to using the Microsoft Management Console (MMC) and command line interfaces to perform Virtual Replicator tasks. The following topics are described in this chapter:

- [Using the Microsoft Management Console Interface](#), page 66
 - [Understanding Snap-ins](#), page 66
 - [Snap-in Panes](#), page 67
 - [Performing Virtual Replicator Tasks](#), page 67
 - [Adding a Snap-in](#), page 68
 - [Managing Remote Computers](#), page 72
- [Using the Command Line Interface](#), page 73
 - [Using Commands to Manage Remote Computers](#), page 73
 - [Abbreviating Commands](#), page 74
 - [Using Uppercase or Lowercase in Commands](#), page 75
- [Using Virtual Replicator with Other Windows Tools](#), page 76
- [Administering Security and Privileges](#), page 77
- [Using Cluster Resources](#), page 78

The sections in this chapter help you get started using Virtual Replicator tools and refer you to other sources of information for performing specific tasks. You can find additional detailed instructions for using Virtual Replicator in the online help available for the MMC and command line interfaces.

Note: Whenever there are two ways to do a task—using Virtual Replicator or using another tool—always use Virtual Replicator. For example, use Virtual Replicator to delete a pool in a cluster, not Cluster Administrator.

Using the Microsoft Management Console Interface

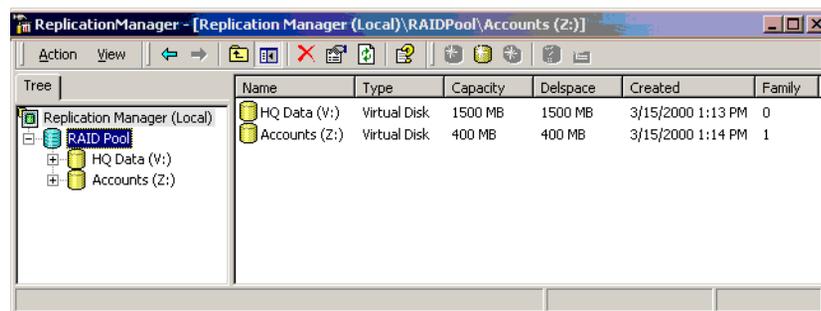
MMC is the management environment from Microsoft, which lets you use different management tools and manage the local and remote computers from within the same window on your desktop. Microsoft and other software vendors are currently migrating their existing management tools into the MMC environment.

Understanding Snap-ins

The management tools that you run in the MMC environment are called **snap-ins**. Virtual Replicator provides one snap-in, Replication Manager.

To start Virtual Replicator using MMC, click **Start > Programs > hp OpenView Storage Virtual Replicator > Replication 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 HP 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 a separate console file that contains all of that person's snap-ins, arranged in a tree view that suits the operator's particular needs.

Console files have the **.msc** extension. Virtual Replicator provides a default console file to get you started quickly, **REPLICATIONMANAGER.MSC**. This file contains a snap-in 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 disks on another computer, add a Replication Manager snap-in and specify the name of the computer you want to manage, as described in [Adding a Snap-in](#) on page 68.

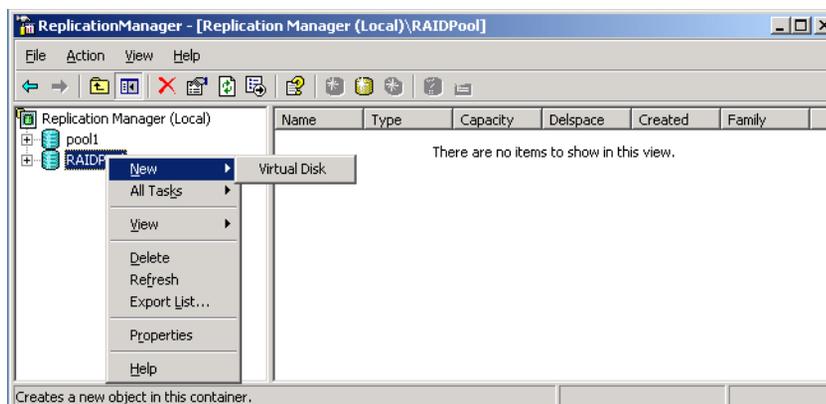
Snap-in Panes

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 Replication 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 Replication 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). |

Performing Virtual Replicator Tasks

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 shortcut menu, as shown in the next screen capture.



Right clicking a pool, virtual disk, or snapshot reveals shortcut menus with different actions for the specific object. The screen above shows the shortcut menu for a pool.

All the actions available on the shortcut menus are also available on the MMC **Action** and **View** menus. For example, in the previous screen, instead of right clicking the pool and clicking **New > Virtual Disk** on the shortcut menu, you could have clicked **Action**, and then **New > Virtual Disk**. These actions are also available as buttons on the toolbar.

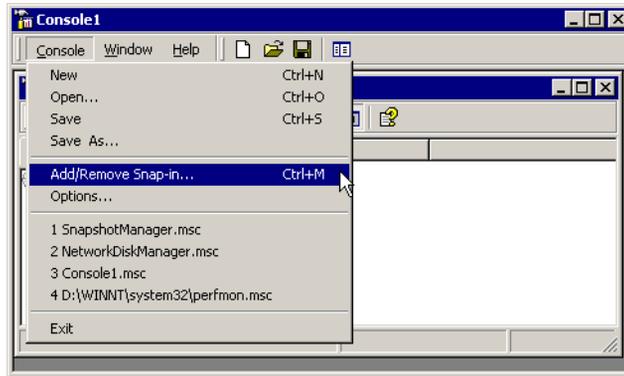
Adding a Snap-in

The default snap-in provided by Virtual Replicator, **Replication Manager**, manages 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 machines.

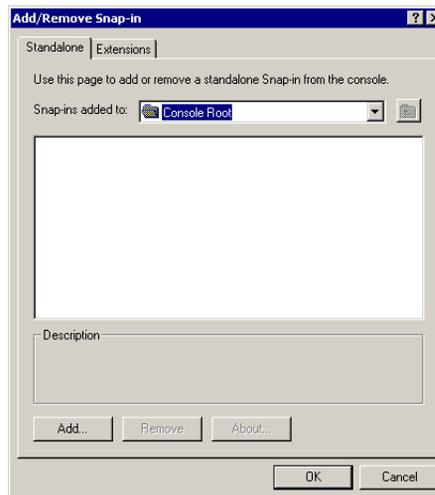
See the MMC documentation for more information about creating a custom console file.

To add a snap-in to your MMC console:

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

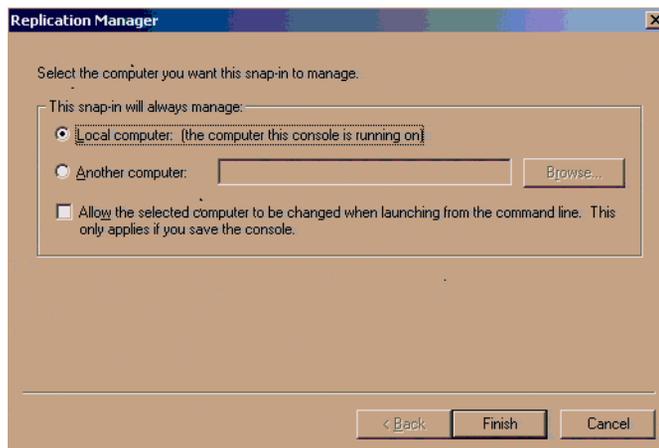


4. On the **Add/Remove Snap-in** screen, use the top list box to select the folder to which you want to add the snap-in. For example, select the **Console Root** folder.



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 **Replication Manager** if you want to manage pools, virtual disks, and snapshots on the specified computer. Click **Add**. The screen for selecting a computer is displayed.



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. Select a computer.
 - 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 Replication Manager snap-in and you type ACCOUNTS1, the snap-in will be labelled Replication 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. Click **OK** to close the **Add/Remove Snap-in** dialog box.
10. 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 Replication Manager.

Managing Remote Computers

You can use Virtual Replicator to manage both local and remote computers and clusters. By default, the Replication Manager snap-in manages the local standalone computer or cluster. Using the snap-in, 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 190 for more information).

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

1. Click **Start > Programs > hp OpenView Storage Virtual Replicator > Replication Manager**.
2. In the left-hand Scope pane, right click the **Replication Manager** snap-in, then click **Connect to another computer**.
3. In Windows 2000 or 2003, 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 to the remote computer and can manage it. 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 68).

Using the Command Line Interface

Virtual Replicator provides a command line interface that lets you manage pools, and the virtual disks and snapshots in them. You can run the SnapMgr commands from any command line. Using the command line interface, you can also create batch jobs to automate tasks. For example, you can write batch jobs to back up your data. The SnapMgr commands are described in Appendix C, [SnapMgr Commands](#).

| SnapMgr Command | Description |
|-----------------|--|
| DRIVES | Shows which drive letters are available to map to virtual disks and snapshots. |
| MANAGE | Controls which computer you manage (only available at the SnapMgr 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. |
| UTILITY | Enables support for incremental backups of virtual disks. |
| VIRTUALDISK | Manages virtual disks. |

The commands have their own command prefix: SnapMgr. For example, this command deletes the snapshot called My Snapshot:

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

The commands also have their own prompt (SNAPMGR). 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 MMC snap-ins, you can use the SnapMgr 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 SnapMgr command prompt or the Windows 2000/2003 prompt.

At the SnapMgr 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 local computer, and then on the remote Accounts computer:

```
SnapMgr> MANAGE /LOCAL
SnapMgr> DRIVES

    Available drive letters: DGHMNPQ

SnapMgr> MANAGE Accounts
SnapMgr> DRIVES

    Available drive letters: GHKMNPQRST
```

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 2000/2003 Command Prompt

If you are at the Windows 2000/2003 command prompt, each command automatically manages the local computer or cluster unless you specify the name of a remote computer after the SNAPMGR command prefix.

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

```
C:\> SNAPMGR DRIVES

    Available drive letters: DGHMNPQ

C:\> SNAPMGR Accounts DRIVES

    Available drive letters: GHKMNPQRST
```

Abbreviating Commands

Abbreviate the commands to three letters or less but do not abbreviate the command prefix (SNAPMGR).

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
```

Using Virtual Replicator with Other Windows Tools

The preceding sections of this chapter provide an overview of the Virtual Replicator tools. This section includes best practices and restrictions for using Virtual Replicator with other Windows tools, such as Disk Management and Cluster Administrator.

Keep in mind the following guidelines:

- Whenever there are two ways to do a task—using Virtual Replicator or using another tool—always use Virtual Replicator. For example:
 - Use Virtual Replicator to delete a pool in a cluster, not Cluster Administrator.
 - Use Virtual Replicator to map drive letters to virtual disks and snapshots, not Disk Management (in Windows 2000/2003).
- Do not use Windows Explorer to monitor free space in a pool; Windows Explorer does not recognize pools.
- Do not use Virtual Replicator at the same time as Windows Disk Management/Logical Disk Manager. Doing so can result in unpredictable behavior.
- Close Logical Disk Manager whenever you are using Virtual Replicator management tools.
- Do not perform Disk Management tasks, such as mapping drive letters or formatting disks, on disks that are part of Virtual Replicator pools.
- Do not perform online volume growth of basic or virtual disks at the same time as other major disk management operations, such as defragmenting and disk checking.
- As an extra precaution, before performing *any* tasks with Disk Management for Windows 2000, use Device Manager to scan for hardware changes in order to update disk information.
 - To open Device Manager, right click on **My Computer**, select **Properties**, and then select the **Hardware** tab.

The recommendations above are general guidelines for using Virtual Replicator. The task descriptions in the Virtual Replicator online help may contain further instructions for using other Windows tools in connection with specific tasks.

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 2000 and Windows Server 2003, 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** on the domain controller. 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.

To use the Virtual Replicator online volume growth feature, you need to have **Full Control** permission at the top level of the volume that you want to grow. Administrators have Full Control access to NTFS volumes on Windows 2000, by default. You do not need to modify this permission unless it has been explicitly denied.

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

When using Virtual Replicator on a Microsoft Cluster Server system, keep the following points in mind:

- When you create pools in a cluster, Virtual Replicator automatically creates all necessary groups and resources. It also sets up correct dependencies, possible owners, and default properties so that your pools can fail over within the cluster.
- For each pool, Virtual Replicator creates a cluster group named **poolname Group** (where **poolname** is the name you gave the pool). It also creates a resource for the pool itself. This resource has the same name as the pool and is of the type **SCE Pool**.
- Do not use Cluster Administrator to set up or rename any of these resources. If you do this, you may not be able to access your data.
- There are no cluster resources for virtual disks or snapshots. The virtual disks and snapshots in a pool automatically appear on whichever node currently owns the pool resource.
- HP recommends that you have a backup copy of the Cluster Registry in the event you need to perform a system recovery.

For example, if you create a pool in a cluster and name it **RAID Pool**, Virtual Replicator creates a new cluster group for the pool called **RAID Pool Group**. This group contains the pool resource, which is also called **RAID Pool**.

By default, all nodes in the cluster are **Possible owners** of the pool resource. After setting up the pool, if you find it necessary to restrict which nodes own the pool resource, you can use Cluster Administrator. See the Microsoft Cluster Administrator documentation for instructions.

After failover, Virtual Replicator cluster groups do not automatically fail back. If you want a group to have failback capabilities, you must use Cluster Administrator to manually configure the group.

See [Pools in a Cluster](#) on page 99 for more information.

Virtual Replicator - A Guided Tour

5

Using a guided tour format, this chapter provides an overview of the Virtual Replicator tools for managing storage pools. The tour shows you how to perform the following tasks:

- [Creating a Pool](#), page 80
- [Creating a Virtual Disk in a Pool](#), page 85
- [Creating a Snapshot of a Virtual Disk](#), page 91

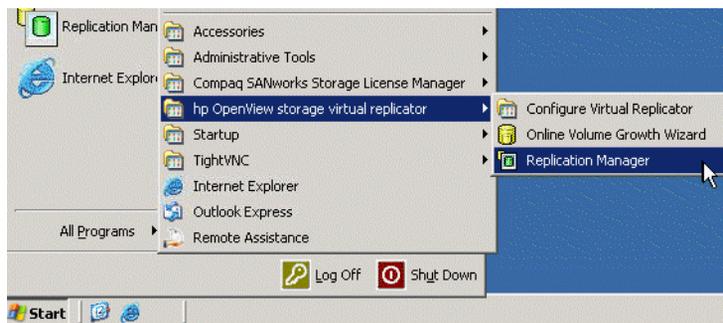
For the purposes of this tour, the tasks occur in the MMC interface. It should take you about 15 minutes to read through the tour—you do not need to follow the steps on your own computer.

Creating a Pool

This part of the tour creates a pool using two disks as the storage units for the pool.

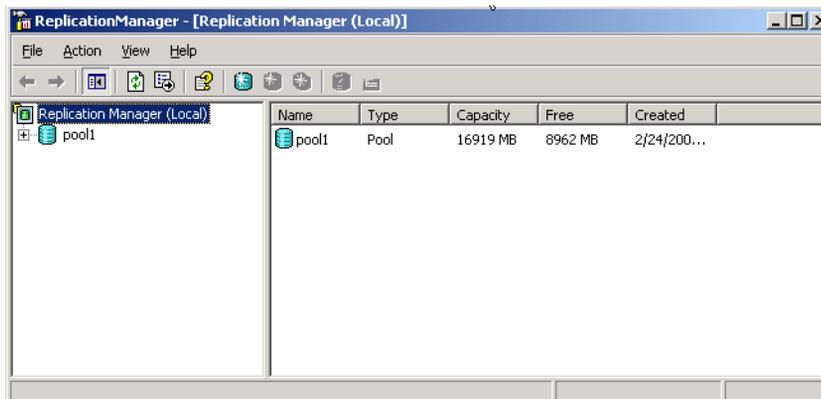
You create pools out of **raw** disks. Therefore, assume that the disks do not contain any partitions.

1. Click **Start > Programs > hp OpenView Storage Virtual Replicator > Replication Manager**.



The MMC splash screen is briefly displayed.

And then the **Replication Manager** console screen is displayed.

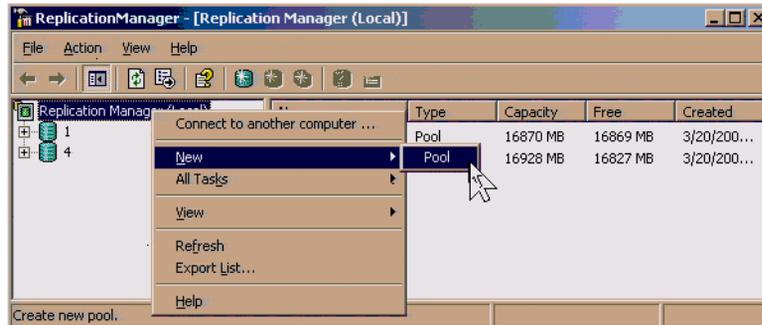


MMC is started and you have opened the **Replication Manager console file**, which helps you get started quickly.

MMC is similar to Windows 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.

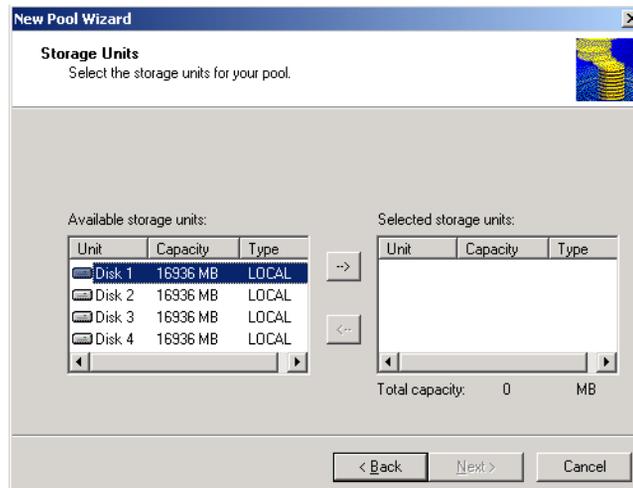
Click the items in the left-hand Scope pane to expand them. The right-hand **Results** pane displays information about whatever you have selected in the left-hand Scope pane.

2. Right click **Replication Manager (Local)**, and then choose **New**, then **Pool**.



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

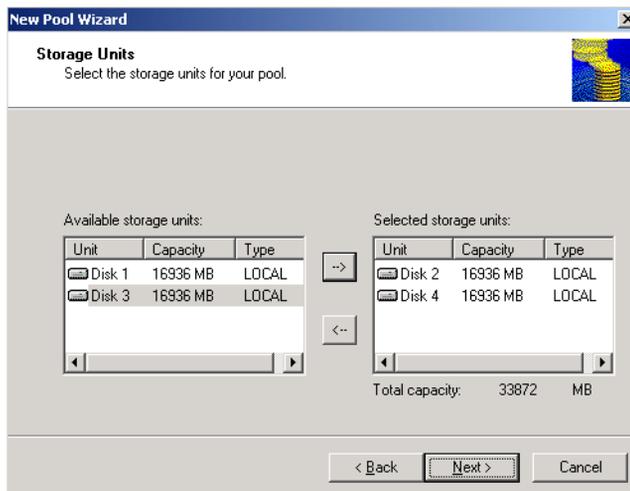
3. Click **Next** on the **New Pool Wizard Welcome** screen. The **New Pool Wizard Storage Units** screen is displayed.



On the left, there is a list of available storage units.

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

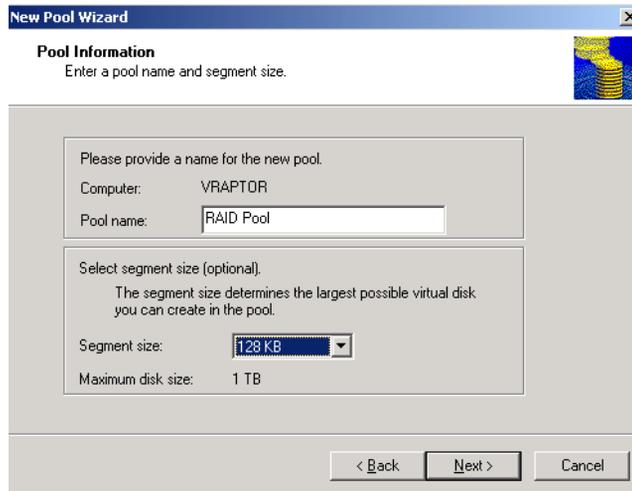
The disks move to the right side of the screen, as shown on the next screen.



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 percent less than this, because some of the disk space is used by internal pool configuration data.

Click **Next**.

5. On the **Pool Information** screen, type the name that you want to give to the new pool. The name can be up to 23 characters and must be different from that of any other pool, virtual disk, or snapshot on the computer or cluster you are managing. This example uses the name **RAIDPool**.



On the same screen, choose the **segment size** or accept the default.

Note: Once you choose the segment size and complete the creation of the pool, you cannot change the segment size.

A **segment** is a fixed-sized unit of contiguous bytes of disk space. The specified segment 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 128 KB limits virtual disks and additional storage units in the pool to a maximum of 1 TB.

Segment size also determines the smallest unit that would be copied out to a snapshot when changes occur on a parent disk. See [Disk Space Usage by Snapshots](#) on page 128 for more information on copy-outs.

Segment size options range from 32 KB to 256 KB. If you choose a large segment size, the maximum size of new virtual disks increases, but the amount of pool space that is needed for snapshots might be greater because of larger copy-outs.

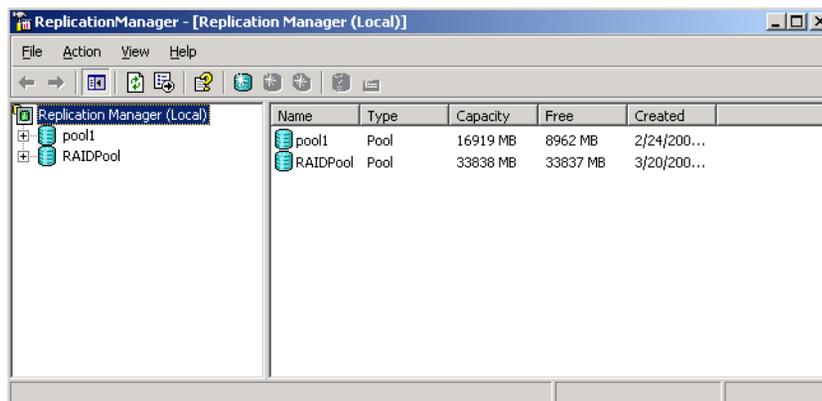
This example uses the default segment size of 128 KB.

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:



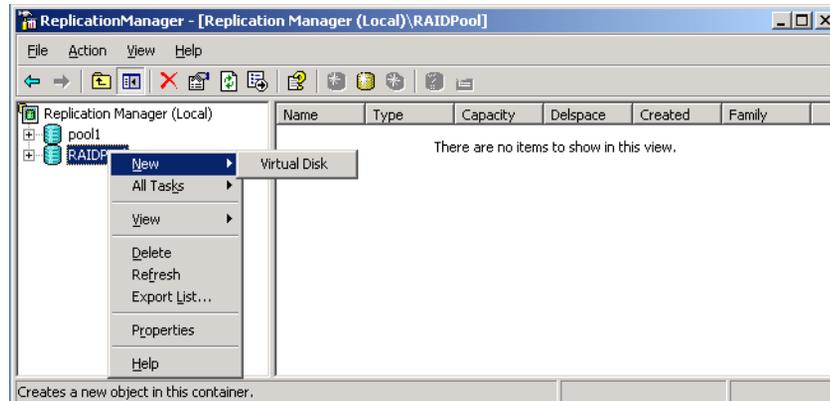
The new pool is listed. Notice that the pool's capacity is 1 MB more than its free space. That is because the capacity is rounded up and the free space is rounded down.

The next part of the tour describes how to create a virtual disk. For more information about pools, see Chapter 6, [Managing Pools](#).

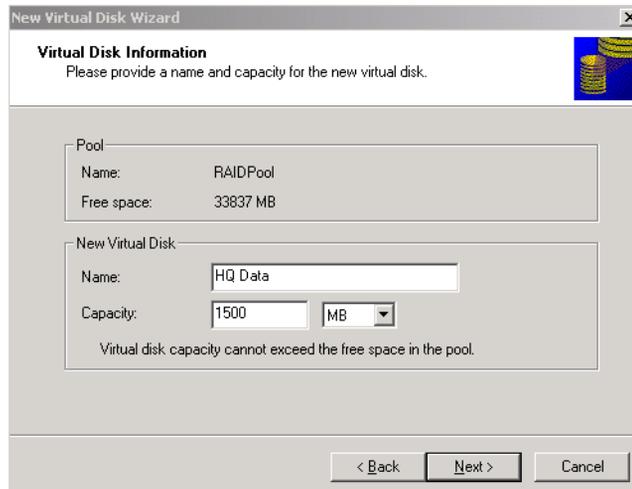
Creating a Virtual Disk in a Pool

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

1. Right click on **RAIDPool**, and then choose **New > Virtual Disk**.



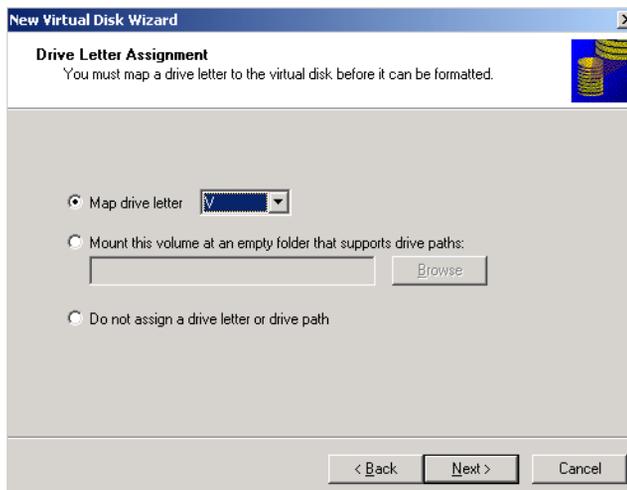
2. The welcome screen of the New Virtual Disk wizard is displayed. Click **Next** to start the wizard.
3. On the next screen, type the name of the new virtual disk and its capacity—this example calls it **HQData** and specifies **1500 MB**—then click **Next**.



4. On the next screen, click either **Map drive letter** or **Mount this volume at an empty folder that supports drive paths**.

To map a drive letter:

- a. Select the drive letter that you want to map to the virtual disk. For example, select **V**.



When mapping drive letters, HP recommends that you select a letter near the end of the alphabet because during system startup:

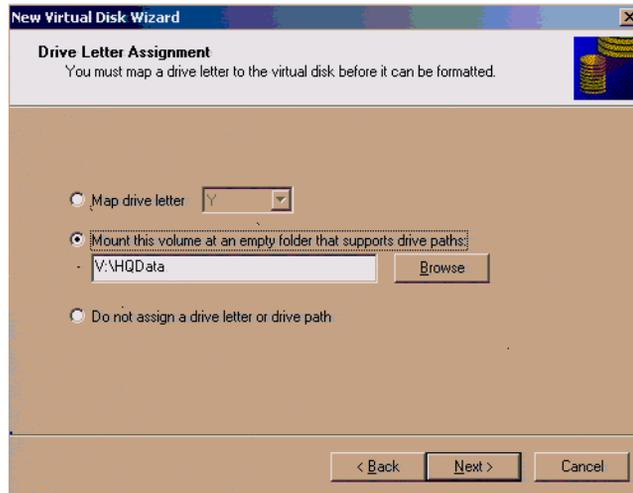
- Drive letters are automatically allocated from the beginning of the alphabet.
- The pool software starts late in the startup sequence.

As a result, the drive letter you choose for the virtual disk might be allocated to another disk that comes online before the pool software starts. Consequently, the virtual disk will have no drive letter assigned. Choosing a letter near the end of the alphabet prevents the allocation of the virtual disk's drive letter to another disk.

- b. Click **Next**.

To create a mount point:

- a. Click **Mount this volume at an empty folder that supports drive paths**.
- b. Click **Browse**.
- c. Select an empty folder on an NTFS volume to mount the drive. For example, select the HQData folder on NTFS volume V: and click **OK**. The mount point path is displayed as V:\HQData.



When creating mount points in a clustered environment, the following best practices are recommended:

- The target disk for a mount point should be a shared disk or a cluster aware disk.
- If the target disk for a mount point is a shared disk, make sure that the mountable folder is seen on both nodes before you create the mount point. This is required for the mount points to failover successfully.
- If the target disk for the mount point is a cluster aware disk, make sure the cluster aware disk fails over before the pool does. This would ensure that the disk is present when pools try rebinding on the other node.
- Nested mount points are not supported. For example, if V:\temp is a mount point, then the mount point V:\temp\xyz is not supported.

- Mount points will be lost when either a cluster upgrade or downgrade is performed. HP recommends manual recreation of mount points after an upgrade to a cluster or a downgrade to a standalone.

d. Click **Next**.

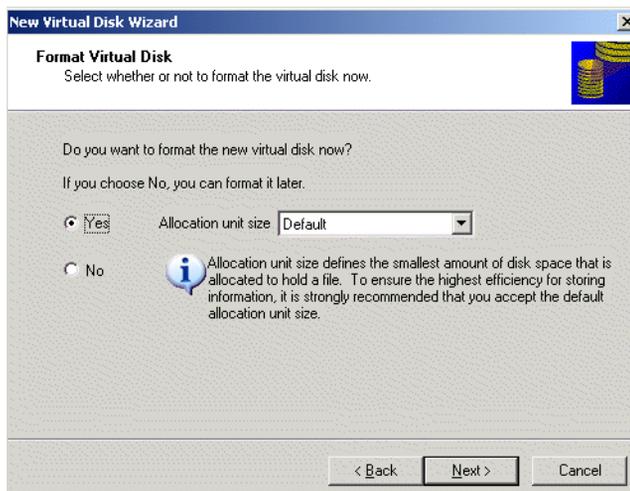
5. If the mount point option was selected, the final wizard screen is displayed. Check that there are no error messages, then click **Finish** to close the wizard.

If the drive letter option was selected, the Format Virtual Disk screen is displayed. Click **Yes** to format the virtual disk. This creates a single partition on the disk, formats it with the NTFS file system, and gives it the same label as the virtual disk name (HQData).

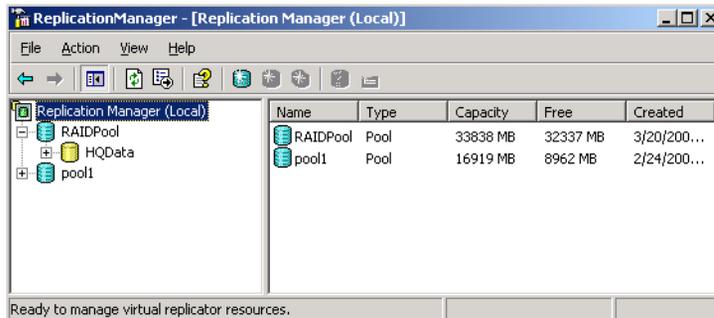
When formatting virtual disks, HP recommends that you accept the default setting for **Allocation unit size**. This parameter defines the smallest amount of disk space that can be allocated to hold a file.

By default, Virtual Replicator sets the allocation unit size to 4096 bytes (4 KB). In most cases, this setting provides the highest efficiency by minimizing the amount of lost space and fragmentation on the volume. Therefore, it is recommended that you not change this value.

Note: You cannot defragment a virtual disk that has an allocation unit size greater than 4 KB.

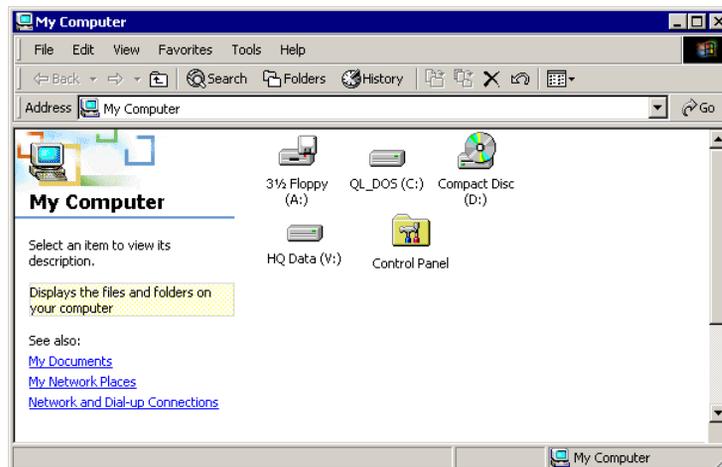


6. Click **Next**.
7. The final wizard screen is displayed. Check that there are no error messages, then click **Finish** to close the wizard.
8. Select the **Replication Manager (Local)** folder.



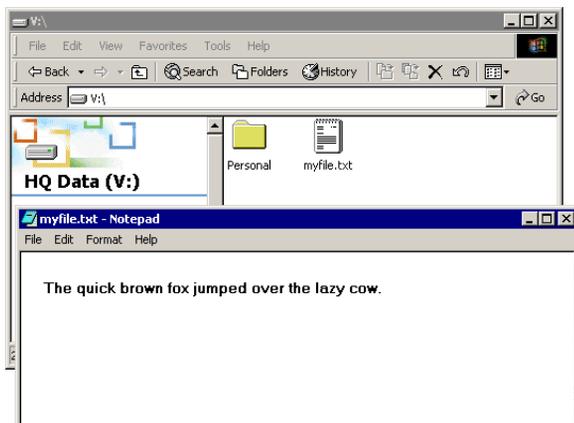
The new virtual disk, **HQData**, is displayed under **RAIDPool**, on the left. The drive letter you mapped to it, **V**, is in brackets. On the right, notice that the pool's free space is now 32337 MB. Free space has dropped by 1500 MB, the capacity of the virtual disk you have just created.

9. From the Windows desktop, open **My Computer**.

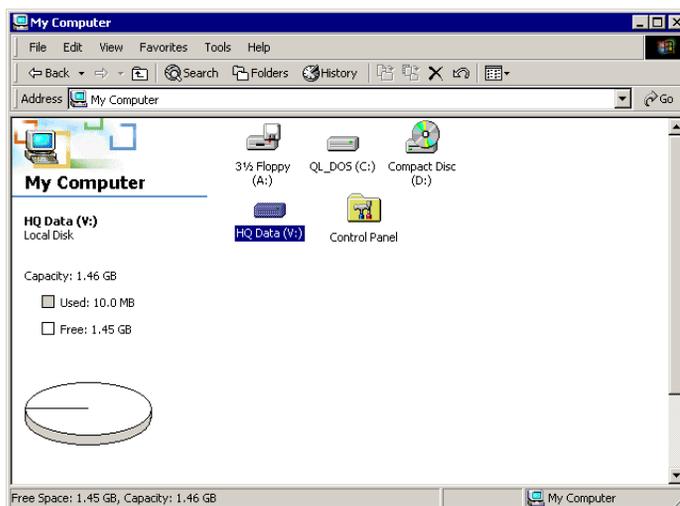


The new disk appears with drive letter **V**, volume label **HQ Data**.

10. Copy a folder called **PERSONAL** and a file called **MYFILE.TXT** to the new disk:



11. 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:



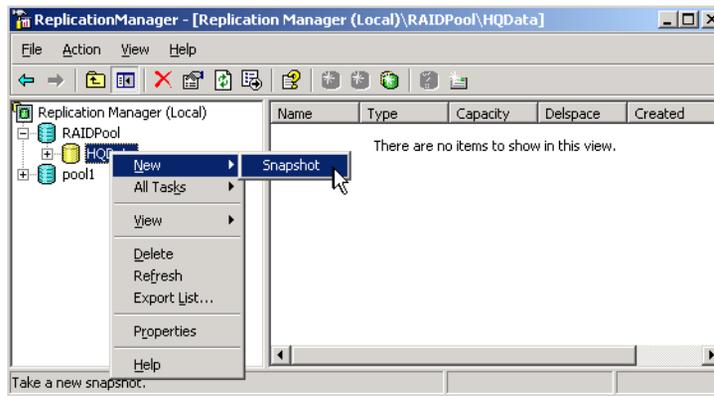
12. For a mounted virtual disk, open Windows Explorer. Go to the V: drive. The new virtual disk is displayed as a mounted drive (V:\HQ Data).

Creating a Snapshot of a Virtual Disk

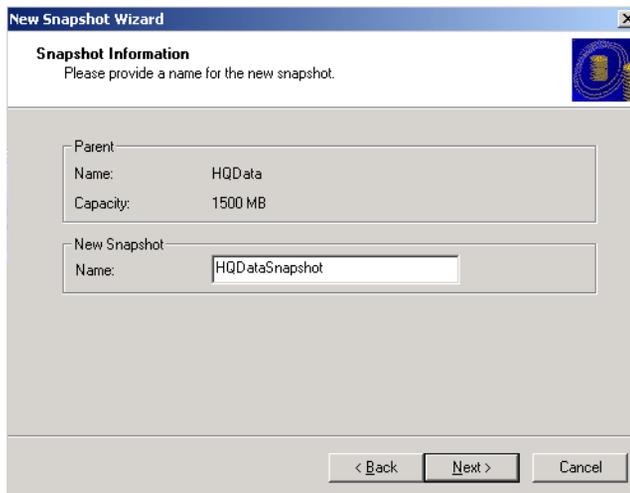
This part of the tour creates a snapshot, called HQ Data Snapshot, of the virtual disk HQData. You will map a drive letter to the snapshot and check that it contains the same files as its parent disk. Then, you will check the free space in the pool to find out how much disk space the snapshot is using.

Note: As a general rule, HP recommends that you not create snapshots while there is rigorous I/O activity occurring on the virtual disk. Doing so could considerably increase the total time needed to create the snapshot.

1. In the **Replication Manager** console, right click the virtual disk that you just created, HQData, then choose **New > Snapshot**.



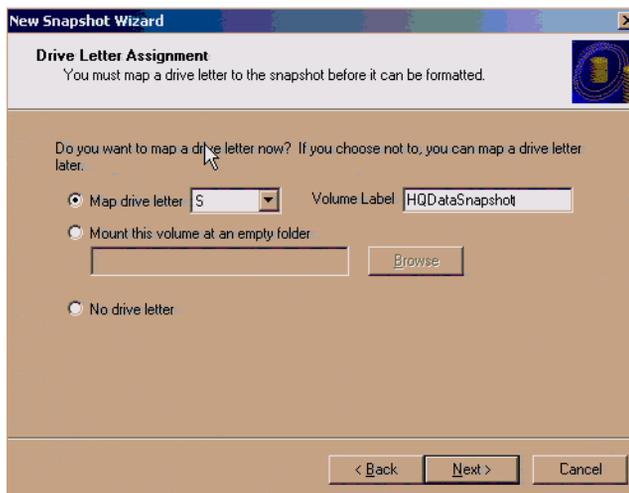
2. Click **Next** on the welcome screen to start the New Snapshot wizard.
3. Type a name for the new snapshot—HQDataSnapshot—then click **Next**.



4. On the next screen, click either **Map drive letter** or **Mount this volume at an empty folder**.

To map a drive letter:

- a. Click **Map drive letter**.
- b. Select the drive letter you want to map to the new snapshot. As with virtual disks, HP recommends selecting a letter near the end of the alphabet. Choose **S** in this example.



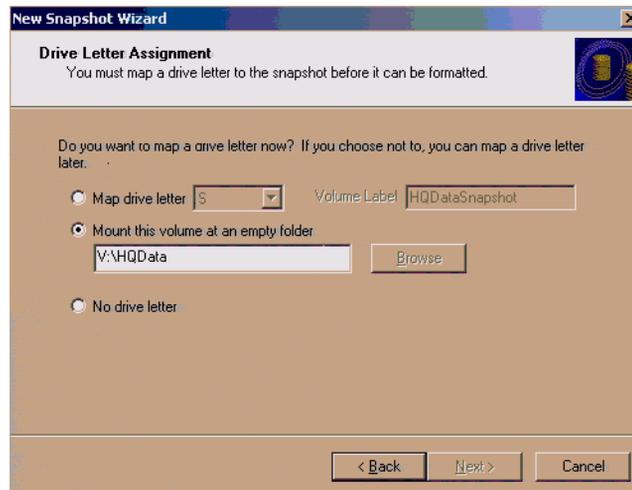
- c. When you assign a drive letter, the new snapshot's volume label is, by default, the same as the parent virtual disk's volume label. That is, multiple disks (a virtual disk and all of its snapshots) have the same volume label. You have the option of changing the name in the **Volume Label** box.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

To avoid confusion between the example virtual disk and snapshot, change the snapshot's volume label to HQDataSnapshot

- d. Click **Next**.

To create a mount point:

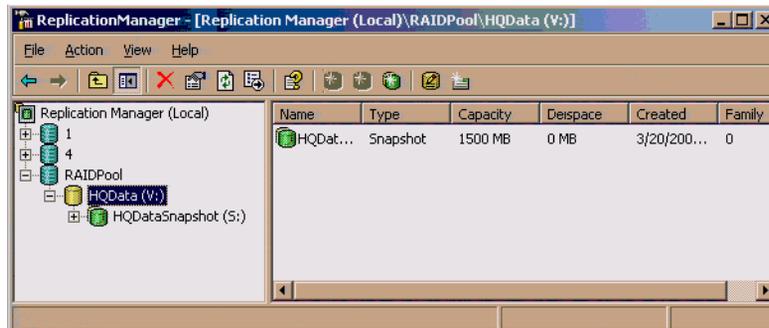
- a. Click **Mount this volume at an empty folder**.
- b. Click **Browse**.
- c. Select an empty folder on an NTFS volume to mount the drive. For example, select the HQData folder on NTFS volume V: and click **OK**. The mount point path is displayed as V:\HQData.



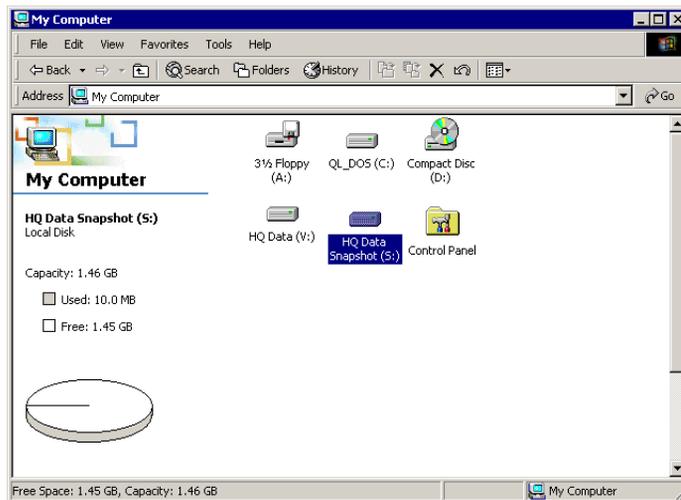
- d. Click **Next**.

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

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

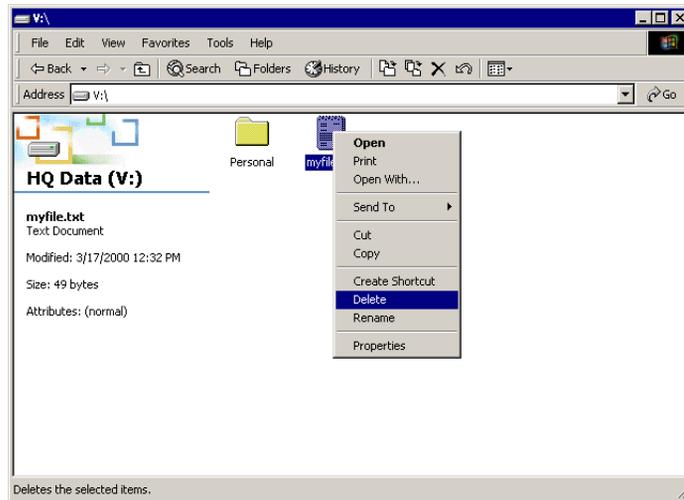


- Open My Computer.

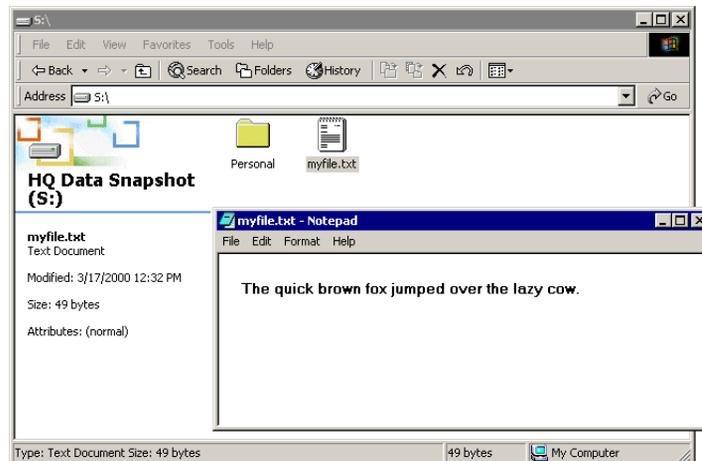


There is the new snapshot, HQDataSnapshot. Select it and look at its free space and capacity on the bottom status bar. It has the same free space as its parent because it contains exactly the same files.

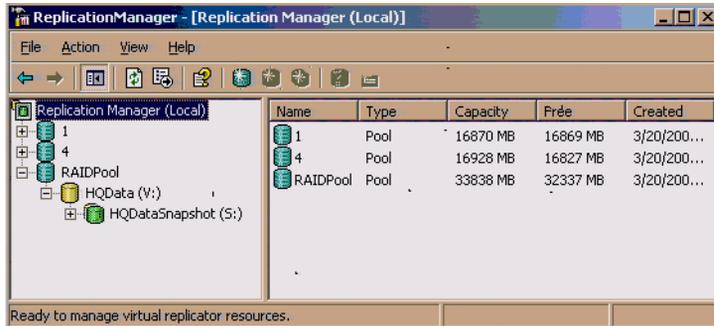
7. You can delete the file Myfile.txt on the parent disk, drive V.



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



8. Finally, go back to the Replication Manager console, and check the pool's free space. It is still 32337 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.



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

Managing Pools

6

This chapter provides information to help you understand and manage pools after you create them. To learn how to create a pool, see the Virtual Replicator online help and also [Creating a Pool](#) on page 80 of this guide.

- [About Pools](#), page 98
- [Pools in a Cluster](#), page 99
- [Additional Storage Units in a Pool](#), page 100
- [About Importing Storage Units](#), page 100
- [About Removing Storage Units from Pools](#), page 102
- [Understanding Pool Free Space](#), page 103
- [Methods for Monitoring Pool Free Space](#), page 103
- [Understanding Delspace](#), page 105
- [Information About Pools](#), page 108
- [About Pool Deletion](#), page 112

About Pools

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.

After you create a pool, Virtual Replicator marks the storage units as offline so that users can no longer access them directly.

See the Virtual Replicator online help and [Creating a Pool](#) on page 80 of this guide for more information on creating pools.

Pools in a Cluster

When you create pools in a cluster, the Virtual Replicator wizard automatically creates a cluster group named *poolname* **Group** (where *poolname* is the name you gave the pool). The wizard then creates a resource of type **SCE Pool** for the pool itself. This resource has the same name as the pool. The wizard brings the pool group and pool resource online.



Caution: Do not use the Microsoft Cluster Administrator to rename the pool resource because you may not be able to access your data.

Although you should not rename the pool resource, it is permissible to rename the pool group. You can have more than one pool in the group, and you can also move pools into different groups.

If you create additional resources in the group for applications or file shares that depend on virtual disks or snapshots in a pool, be sure to set cluster dependency on the pool. Setting dependency on the pool results in the pool starting up before the application or file share starts up.

Specifying Possible Owners for Pool Resources

When you create pools in a cluster, all nodes in the cluster can, by default, host the pool resource. That is, all nodes are **Possible owners** of the resource. If you find it necessary to restrict which nodes own the pool resource, use Cluster Administrator.

In Cluster Administrator, the pool resource **Properties** page, **General** tab lists the nodes that are **Possible owners**. You can add and remove nodes using this page. See the Microsoft Cluster Administrator documentation for complete instructions.

If you remove a node from the list of possible owners, the pool resource and group will not be able to fail over to that node. If you want the group to have failover capabilities, be sure to specify more than one possible owner.

Additional Storage Units in a Pool

After you create a pool, you can add storage units to it at any time, even as users access its virtual disks and snapshots. Storage units must meet the following criteria:

- You can have up to eight 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 logical storage unit that can be included in a pool is 1 TB.
- Storage units should all have the same read-write, redundancy, and failure characteristics. For example, the units should all be standard disks, or they should all be RAID 5 storage sets, 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.

After you add storage units, they are automatically marked as offline so that users can no longer access them directly.

About Importing Storage Units

When creating pools with earlier versions of Virtual Replicator, you could not add storage units that contained data partitions. Storage units were required to be raw and unformatted. If you wanted to use a partitioned disk in a pool, you would have to back up all live data onto a secondary storage, delete the partitions on the disks, and then add the disk into the pool.

This version of Virtual Replicator allows you to import disks with live data directly into a pool with no disruption or downtime. The Import Unit feature makes it possible for you to quickly migrate existing physical storage to your virtualized storage environment.

When you import a partitioned storage unit, Virtual Replicator automatically creates a virtual volume for each partition in the unit. The applications that used data on the physical unit can continue seamlessly to access the data on the virtual disk.

If the physical unit has unused capacity, it will be added to the pool's free space after the import operation. You can then use this additional capacity for other data storage purposes.

The following restrictions apply when you import a storage unit:

- Before importing a storage unit, make sure the pool has free space in an amount equal to or greater than 1% (approximately) of the storage unit's capacity.
- Virtual Replicator does not support the FAT file system. Therefore, if the storage unit you want to import has FAT partitions, you should convert them to NTFS before you import the unit.
- In a cluster, the storage unit to be imported must be a shared disk and must also be a cluster resource.
- You can add a single drive to a RAID set or a LUN that is committed to a pool. However, you will not be able to use the extra space that the drive would add. To use the extra space, you must add the drive to the pool as a separate LUN.

You can, however, add a single drive to the storage system as long as it does not interfere with the storage units that are committed to the pool.

See the Virtual Replicator online help for complete instructions on importing storage units.

About Removing Storage Units from Pools

The Remove Unit feature of Virtual Replicator lets you remove storage units from a pool if the storage is needed for other purposes. You can remove units only if no data has been written to them.

You cannot remove a storage unit from a pool if the unit is being used by virtual disks. Since it is not possible to identify which storage units are being used by particular virtual disks, the best time to use the Remove Unit feature is immediately after you add a storage unit to a pool.

See the Virtual Replicator online help for complete instructions on removing storage units.

Understanding Pool Free Space

The free space in a pool is disk space that is not currently allocated to virtual disks, snapshots, or internal configuration data. You should take care to not let pools run out of free disk space. If a pool has no free space, further writes to it may fail, and you could lose data. You can 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 disk space 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, it could become full. 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, do not try to free up space by deleting files. Deleting files causes segments of data to be copied out to snapshots, which uses up even more free space in the pool and can cause additional data loss.

To manage pools effectively, you should regularly monitor their free space and understand how much space would be freed up by deleting virtual disks and snapshots. These topics are discussed in the following sections:

- [Methods for Monitoring Pool Free Space](#) on page 103
- [Understanding Delspace](#) on page 105

Methods for Monitoring Pool Free Space

Use any of the following methods to monitor free space:

- Visually monitor pool free space in the Replication Manager MMC window, as described in [Using Replication Manager](#).
- Check the System Event Log for Virtual Replicator warnings and errors about pool free space.

When the free space falls to less than 30 percent, 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 percent. These records have **cpqvrbus** in the source column (on Windows 2000).

See [Virtual Replicator Driver Events](#) on page 165 for a full description of these messages.

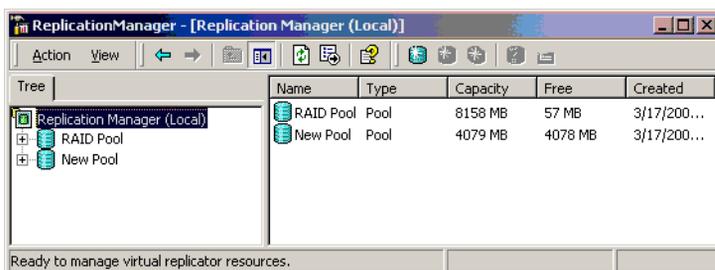
- If you want to gain further control over the monitoring of pool free space, use the Virtual Replicator Lifeguard service, as described in Appendix B, [Virtual Replicator Lifeguard Service](#).

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.

Using Replication Manager

Below is an example of monitoring a pool's free space in the **Replication Manager** MMC window.

In this example, **RAID Pool** has a capacity of **8158 MB**, but with 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.

Because Windows Explorer does not recognize pools, it 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 earlier in the **Replication Manager** window, is only 57 MB.

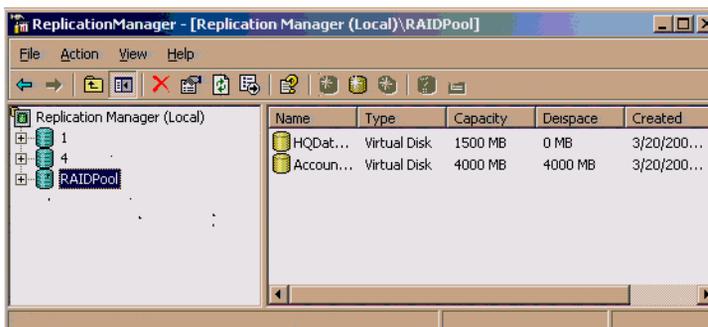
Likewise, deleting files has different effects in Windows Explorer and Replication Manager. For example, if you deleted files on drive Z: (VirtualDisk), you would see the free space rising in Windows Explorer. However, in Replication Manager, free space would be unchanged, because deleting files does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

Understanding Delspace

The **Delspace** of a virtual disk or snapshot is the amount of space you gain in the pool if you delete the virtual disk or snapshot. The free space in the pool increases by the Delspace amount when you delete that virtual disk or snapshot.

Virtual Disk Delspace

To display the Delspace of all the virtual disks in a pool, select the pool in the left-hand Scope pane of the Replication 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 any snapshots in its family, its Delspace is 0. That is because you cannot delete the virtual disk to free up pool space when there are snapshots present. As soon as you delete the last snapshot, you are able to delete the virtual disk, and the Delspace of the virtual disk changes from 0 to its capacity.

Snapshot Delspace

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 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.

Delspace of Single and Multiple Snapshots

- When there is only one snapshot in a family, its Delspace is actually the amount of space that it is using in the pool.
- 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.
- When you delete two snapshots in the same family, you may free up more space than the sum of their Delspaces.

In the example on page 106, if 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 is now used exclusively by Accounts Tue.

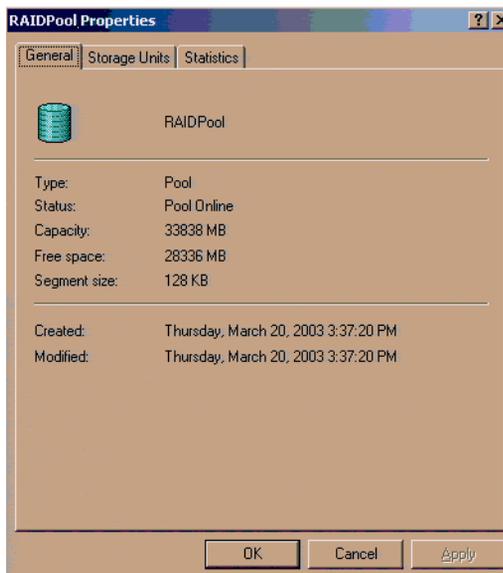
Understanding Delspace and closely monitoring the free space in a pool will help you to prevent pool-full situations. You can gain more knowledge about pools by viewing the abundant information available for each pool, as described in the next section.

Information About Pools

The Replication Manager window shows summary information about all the pools on a standalone 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. |

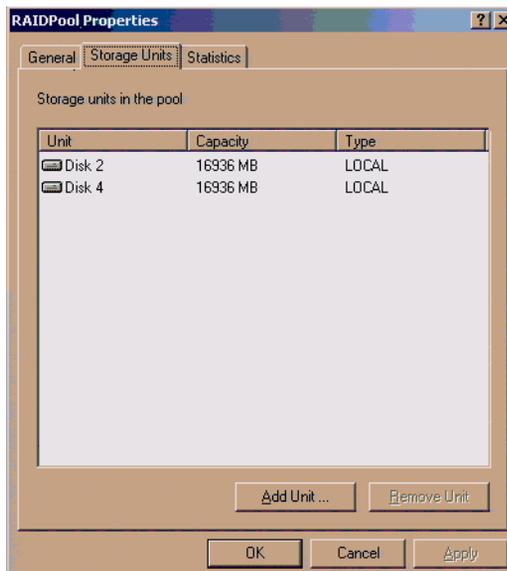
You can learn more specific information by right-clicking the pool and opening the **Properties** page. The page contains three tabs: **General**, **Storage Units**, and **Statistics**.



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 the pool was last modified. This is the date when a storage unit was last added to it.

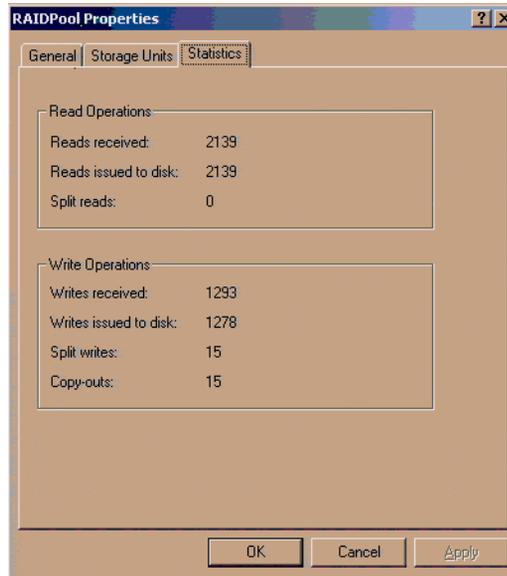
The **Storage Units** tab shows the storage units that are in the pool.



For each unit, the page shows the following information:

- Its disk number. In a cluster, the disk number is 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.



The tab shows the following counters, which are reset to zero when the computer restarts, or, in a cluster, when the pool fails over.

| 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 number 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 number 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. |

About Pool Deletion

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

Deleting the pool frees its storage units and brings them online so that you can access them directly. You can use the storage units 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.

Pool Deletion in a Cluster

Deleting a pool automatically deletes the pool resource (of type **SCE Pool**). When the pool is deleted, the pool group is also deleted, if the group contains no other resources.

Managing Virtual Disks



After you create pools (described in [Managing Pools](#)), you can 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 provides information to help you understand and manage virtual disks after you create them. To learn how to create a virtual disk, see the Virtual Replicator online help and also [Creating a Virtual Disk in a Pool](#) on page 85 of this guide.

The following topics are discussed in this chapter:

- [About Virtual Disks](#), page 114
- [About Drive Letter Mapping and Unmapping](#), page 115
- [Virtual Disk Formatting](#), page 116
- [Virtual Disk Defragmenting](#), page 117
- [Virtual Disk Properties](#), page 118
- [Methods for Restoring Virtual Disks](#), page 121
- [About Deleting Virtual Disks](#), page 124
- [About Growing Virtual Disks](#), page 125

Note: Whenever there are two ways to do a task—using Replication Manager or using another tool—always use Replication Manager. For example, use Replication Manager to map drive letters to virtual disks, not Disk Administrator.

About Virtual Disks

After you create a pool, you can create up to eight virtual disks in the pool. When you create a virtual disk in a pool, you do not specify which storage units it uses. The virtual disk may use disk space from anywhere in the pool.

You can set the capacity of each virtual disk to be at least 10 MB, and at most 2 TB. The capacity is limited by the pool's free space, its segment size, and by the policies in effect for that pool.

Note: Although you can create 2-TB virtual disks, current firmware and hardware support only four 2-TB virtual disks in one pool.

See the Virtual Replicator online help for full instructions on creating virtual disks.

Virtual Disks in a Cluster

There are no cluster resources for virtual disks as there are for pools. The virtual disks in a pool automatically appear on whichever node currently owns the pool resource.

About Drive Letter Mapping and Unmapping

Use Virtual Replicator, not Disk Administrator or Disk Management, to map or unmap drive letters for virtual disks or snapshots. See the Virtual Replicator online help for full instructions on mapping and unmapping.

When mapping drive letters, HP recommends that you select a letter near the end of the alphabet because, during system startup:

- Drive letters are automatically allocated from the beginning of the alphabet.
- The pool software starts late in the startup sequence.

As a result, the drive letter you choose for the virtual disk or snapshot might be allocated to another disk that comes online before the pool software starts. Consequently, the virtual disk or snapshot will have no drive letter assigned. Choosing a letter near the end of the alphabet prevents the allocation of the drive letter to another disk.

Drive Letter Mapping in a Cluster

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

Virtual Disk Formatting

When a virtual disk is created, it consists of a single partition that consumes the entire disk. You can use Virtual Replicator to format the virtual disk to, for example, wipe the disk clean by erasing all of its data.

When you format a virtual disk, you can change NTFS disk parameters, such as allocation unit size. This parameter defines the smallest amount of disk space that can be allocated to hold a file. A default allocation unit size of 4 KB is set by Virtual Replicator. This default setting minimizes the amount of lost space and the amount of fragmentation on the volume. Therefore, in most cases, it is best to use the default allocation unit size.

Another setting that can be changed during formatting is the virtual disk's volume label. This is the name used by Windows to identify the virtual disk. Changing the volume label does not change the name of the virtual disk in Virtual Replicator.

You can also use Disk Administrator or Disk Management 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.



Caution: Never use Disk Administrator or Disk Management to repartition a virtual disk, because doing so could cause data loss.

Virtual disks must have mapped drive letters before they can be formatted. See the Virtual Replicator online help for full instructions on formatting virtual disks.

Virtual Disk Defragmenting

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

If you defragment 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 movement causes unnecessary copy-out operations that degrade performance and waste disk space. To avoid unnecessary copy-out operations, delete all the snapshots on a virtual disk before defragmenting. See [About Snapshots](#) on page 128 for more information.

Note: Because of a Windows restriction, you cannot defragment a virtual disk that has an allocation unit size greater than 4 KB. Accordingly, Virtual Replicator uses 4 KB as a default allocation unit size. See [Virtual Disk Formatting](#) on page 116 for more information.

Virtual Disk Properties

Information about virtual disk properties is available from both the Virtual Replicator Replication Manager and the Windows **My Computer** icon.

Using Virtual Replicator

You can view information about a virtual disk by selecting the virtual disk's pool in the left-hand Scope pane of the Replication 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. |

More information is available by opening the virtual disk's **Properties** page.

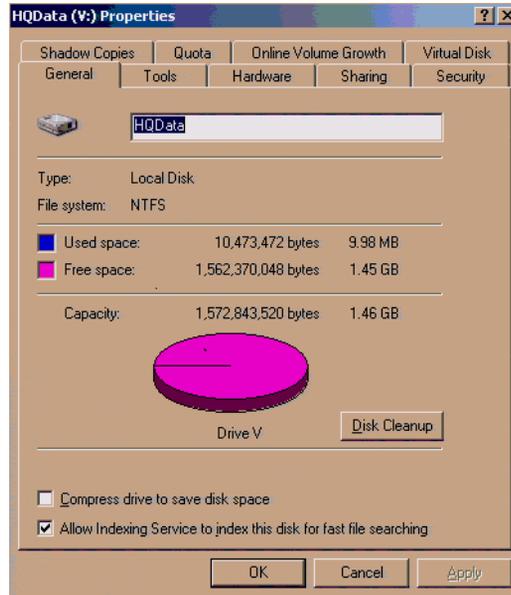


The **General** tab of the **Properties** page shows the same information for the virtual disk as appears in the MMC pane: name, type, capacity, delspace, family, and date created. It also gives the pool, disk number, and drive letter of the virtual disk.

The **Scheduled Tasks** tab lists any tasks that have been scheduled for this virtual disk. See [Task Scheduling Wizards Overview](#) on page 152 for more information.

Using My Computer

Virtual disks appear as icons in the My Computer folder. You can open a Windows **Properties** page for each virtual disk, as shown in the next screen.



The **General** tab shows the following information about the virtual disk, such as its disk number. The disk number may change when you restart the computer or, in a cluster, when the pool resource fails over.

| 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. |
| 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 118).

Methods for Restoring Virtual Disks

When data on a virtual disk becomes corrupted, you can restore the virtual disk from your normal tape backup. However, Virtual Replicator offers several methods for using snapshots to restore virtual disks that can help to minimize downtime.

- SnapBack can be used to automatically restore virtual disks (see [The SnapBack Method of Restoring Virtual Disks](#)).
- Other methods can be used to manually restore virtual disks (see [Other Methods for Restoring Virtual Disks](#)).

The SnapBack Method of Restoring Virtual Disks

Virtual Replicator's SnapBack feature efficiently restores production volumes from existing snapshots.

SnapBack uses the contents of a snapshot to restore the contents of the parent virtual disk. You can use the snapshot to recover data without performing a file copy to an intermediate location and consuming disk space in the process.

The SnapBack process is automatic. That is, there is no need for you to map and remap drives, copy files, or delete snapshots or virtual disks. SnapBack performs all of these steps for you.

As part of the process, the source snapshot is retained, but all snapshots that are older than the source for the SnapBack are deleted. Deleting earlier snapshots ensures:

- Integrity of the data on the restored virtual disk.
- Consistency between the virtual disk and subsequent snapshots.

When SnapBack is performed, Virtual Replicator deletes all data from the virtual disk and replaces it with the snapshot's data. If the virtual disk changed between the time you made the snapshot and the time you perform SnapBack, the new data will not be retained.

The duration of the SnapBack process can vary, depending on several factors, such as the amount of data on the snapshot, the number of storage units in the pool, and your hardware configuration. See [Example of Using SnapBack](#) for more information.

Example of Using SnapBack

The following example illustrates one possible implementation of SnapBack.

The event that triggers a need for SnapBack is a virtual disk that becomes corrupt during software testing. The scenario below shows the results of SnapBack in a specific hardware and software environment. Other configurations will experience SnapBack differently, although the successful conclusion with a restored virtual disk will be the same.

| SnapBack Scenario | |
|----------------------------|---|
| Configuration | HP StorageWorks Proliant system RA4200 SCSI Controller External storage box containing (14) 18.2-GB drives 512 MB memory (2) 500-MHz processors |
| Virtual Replicator Setting | Pool: 242 GB capacity 1 virtual disk (vdisk_1): 122 GB capacity; 60 GB data written 1 snapshot (snap_1) taken of the virtual disk |
| Event | Install new software for testing purposes on vdisk_1, writing an additional 60 GB of data. (snap_1's Delspace increases to 60 GB.) The software corrupts the data on vdisk_1. |
| Response | Start Virtual Replicator, select snap_1, and run SnapBack. vdisk_1 is restored with the 60 GB of data from snap_1. The operation takes five hours. Note: Virtual Replicator does not restore the 60 GB of data that was written to vdisk_1 <i>after</i> the snapshot was created. This data is lost. |

Other Methods for Restoring Virtual Disks

SnapBack is the most direct way to restore a virtual disk. However, there are three other methods by which you can use snapshots to manually perform an online restore of a virtual disk:

- **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. This method is the manual equivalent of SnapBack.

Although Method 1 is the quickest of these manual techniques to make data available to users, it leaves the original virtual disk as an orphan. This method also requires more pool space and reduces performance because of additional copy-outs.

Method 2 is best used for restoring individual files that have been deleted or corrupted.

You can use Method 3 to restore an entire volume. However, this method causes temporary space consumption, because of the presence of two virtual disks and the need to preallocate pool space. To conserve disk and pool space, SnapBack is the preferred method.

About Deleting Virtual Disks

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

The following guidelines apply to deleting virtual disks:

- Deleting a virtual disk destroys all the data stored on the disk. If the virtual disk contains any data you want to keep, back up the data.
- You cannot delete a virtual disk that has snapshots.
- After you delete a virtual disk, use the Windows Scheduled Tasks applet to delete any scheduled tasks for the disk. Otherwise, if you reuse the previous virtual disk name for another virtual disk, the old tasks will run on the new disk.

About 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. See Chapter 9, [Online Volume Growth](#), for detailed information about this feature.

Managing Snapshots

8

After you create a virtual disk (described in [Managing Virtual Disks](#)), you can begin to take snapshots of it. A snapshot is an instant replica of a virtual disk. Snapshots are useful when you want to test new applications or when you need to restore data after a disk error or failure.

This chapter provides information to help you understand and manage snapshots after you create them. To learn how to create a snapshot, see the Virtual Replicator online help and also [Creating a Snapshot of a Virtual Disk](#) on page 91 of this guide.

The following topics are discussed in this chapter:

- [About Snapshots](#), page 128
 - [Disk Space Usage by Snapshots](#), page 128
 - [Best Practices for Using Snapshots](#), page 130
 - [Snapshots in a Cluster](#), page 130
- [Automatic Snapshots](#), page 131
- [Snapshot Drive Letters](#), page 132
- [Snapshots for Backups](#), page 133
- [Snapshot Properties](#), page 134
- [About Deleting Snapshots](#), page 136

Note: Whenever there are two ways to do a task—using Replication Manager or using another tool—always use Replication Manager. For example, use Replication Manager to map drive letters to snapshots, not Disk Administrator.

About 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 a snapshot of either a virtual disk or another snapshot. The original virtual disk, all of its snapshots, and all the snapshots of those snapshots are known as a **family**. You can have up to 12 snapshots in each 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.

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.

Disk Space Usage by Snapshots

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**.

When a copy-out occurs, the snapshot starts using disk space in units called **segments**. A segment is a fixed-sized unit of contiguous bytes of disk space. Segment size is one of the parameters you set when you create a pool (see page 80). 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 [Figure 2](#).

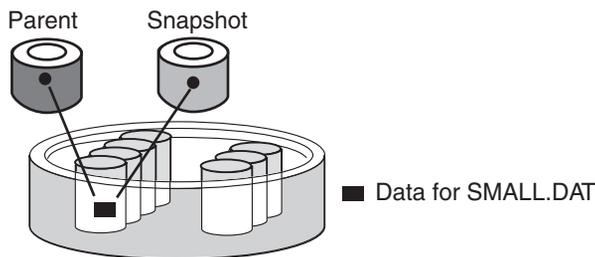


Figure 2: Example of a New Snapshot's Disk Space Usage

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.

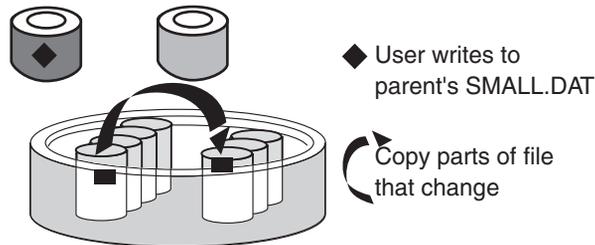


Figure 3: A Copy-out Operation

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 can be copied in one segment of 128 KB. But for a normal file, the parts that change would be copied in multiple segments.

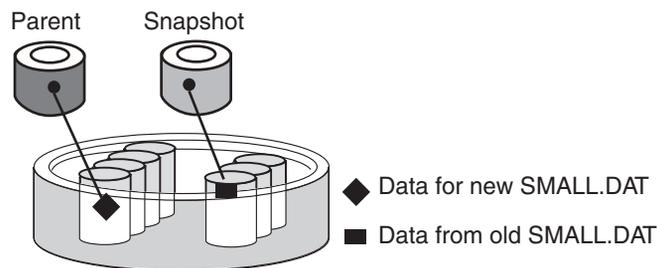


Figure 4: Disk Space Usage after a Copy-Out

When a virtual disk segment has been copied-out, subsequent writes to the segment do not incur copy-out operations. The snapshot has its own copy of the original segment, which does not change. Therefore, there is no need to copy it again.

Best Practices for Using Snapshots

To maintain optimum system performance, HP recommends that you avoid keeping snapshots for long periods of time if data is changing frequently.

When data on a parent disk changes, it causes segments to be copied-out to each of the parent disk's current snapshots. Virtual Replicator keeps track of copy-outs by associating an address with each segment of data. During startup, Virtual Replicator pools "rebind" by re-establishing these address connections. The rebinding process is usually quick, but can take much longer if you have many snapshots containing large amounts of changed data (copy-outs).

Accordingly, you should not consider Virtual Replicator snapshots your only means of backup. HP strongly recommends that you use Virtual Replicator in conjunction with a standard backup tool, and that you use the following guidelines:

- Avoid keeping snapshots for long periods of time if there are large changes in data.
- Use standard backup tools to back up your snapshots.
- Delete snapshots promptly after backing them up.

See [Snapshots for Backups](#) on page 133 for more information on the best practice for using snapshots to back up virtual disks. See [Rules for Using Virtual Disks and Snapshots](#) on page 56 for other guidelines.

HP recommends that you not create snapshots while there is rigorous I/O activity occurring on the virtual disk. Doing so could considerably increase the total time needed to create the snapshot.

Snapshots in a Cluster

Virtual Replicator does not create any cluster resources for snapshots as it does for pools. The snapshots in a pool automatically appear on whichever node currently owns the pool resource.

Automatic Snapshots

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 10, [Scheduling 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 D, [SmartSnap Command](#).

Snapshot Drive Letters

The procedures for mapping and unmapping drive letters to and from snapshots are the same as those for virtual disks. See [About Drive Letter Mapping and Unmapping](#) on page 115 for more information.

Snapshots for Backups

The basic steps involved in using snapshots to back up virtual disks include:

- Stopping an application.
- Creating a snapshot.
- Restarting the application
- Executing the backup using your backup tool.
- Deleting a pre-existing snapshot.

You can use the SnapMgr commands to carry out these steps and to create batch jobs to automate your backups. Alternatively, you can use the Snapshot for Backup wizard (see page 152). This wizard also performs an e-mail or pager notification when the backup is complete.

See the Virtual Replicator online help for detailed descriptions of these backup tasks.

Snapshots for Restoring 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 storage website contains information on using snapshot technology for backing up and restoring in specific application environments, such as Microsoft Exchange. To learn more about these solutions, go to:

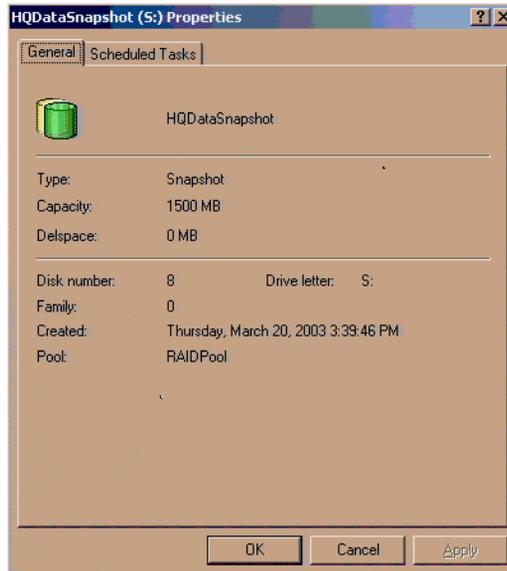
<http://h18000.www1.hp.com/products/sanworks/vr/index.html>.

Snapshot Properties

Information about a snapshot's properties is available by selecting the snapshot's virtual disk in the left-hand pane of the Replication 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 value is used by the file system and all other software on the system. For example, Windows Explorer uses the capacity to display the snapshot's size. Capacity also represents 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. Delspace is the amount of pool space used exclusively by this snapshot—the space that the snapshot 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, right click the snapshot name, then choose **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 [Task Scheduling Wizards Overview](#) on page 152 for more information.

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

After you delete a snapshot, use the Windows Scheduled Tasks applet to delete any scheduled tasks for the snapshot. Otherwise, if you reuse the previous snapshot name for another snapshot, the old tasks will run on the new snapshot.

See Chapter 10, [Scheduling Tasks](#), and Appendix A, [Virtual Replicator Policy Management](#), for additional methods for deleting snapshots.

Online Volume Growth

9

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 provides information and examples to help you understand and prepare for online volume growth. The following topics are discussed:

- [Online Volume Growth for Basic and Virtual Disks](#), page 138
 - [Understanding Basic, Virtual, and Dynamic Disks](#), page 138
 - [Online Volume Growth Planning](#), page 140
- [About Growing Virtual Disks](#), page 142
- [About Growing Basic Disks](#), page 143
- [Online Volume Growth in Windows Explorer](#), page 150

The Virtual Replicator online help contains detailed instructions for performing online volume growth.

Online Volume Growth for Basic and Virtual Disks

Virtual Replicator 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 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 MS-DOS, Windows 3.1, Windows 95/98, and Windows NT, and have been used extensively in the past, basic disks 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 an HP 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.

Although it is possible to convert a basic disk to a dynamic disk while retaining volumes on the disk, it is not possible to revert back to basic 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.**

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

Online Volume Growth Planning

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. The larger virtual disk 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 HP 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 HP Array Configuration Utility.

- Verify volume configuration compatibility.

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, mirrorsets, stripesets, and stripesets 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.

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.

- Check for correct security settings.

To use the online volume growth feature, you need to have **Full Control** permission at the top level of the volume that you want to grow. By default, administrators have Full Control access to NTFS volumes on Windows 2000 and Windows Server 2003. You do not need to modify this permission unless it has been explicitly denied.

About Growing Virtual Disks

A pool must have available free space before you can grow one of its virtual disks. If the pool doesn't have adequate free space, 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 eight 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 [Adjacent Space on Basic Disks](#) on page 144. The only way to grow a pool is by adding storage units to it.

After you add storage units to create pools, you can create virtual disks. You can also grow a virtual disk up to a maximum size of 2 TB, 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.

You can grow a virtual disk in any of the following ways:

- Selecting the disk in the Replication Manager MMC interface.
- Selecting the disk in Windows Explorer. See [Online Volume Growth in Windows Explorer](#), on page 150 for more information.
- Using the Online Volume Growth wizard on the Virtual Replicator Program menu.
- Using the command line. See [Grow a Virtual Disk](#) on page 227 for more information.

See the Virtual Replicator online help for detailed instructions on 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.

About Growing Basic Disks

You can grow a basic disk in any of the following ways, by:

- Using the Online Volume Growth wizard on the Virtual Replicator Program menu.
- Selecting the disk in Windows Explorer.

The rest of this section discusses the steps involved in understanding and preparing for online volume growth of basic disks. See the Virtual Replicator online help for detailed instructions on growing basic disks.

Basic Disk Preparation

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.

Only partitions and free space adjacent to the end of a particular volume are eligible to be incorporated into the volume during volume growth. The diagram below shows a configuration that allows the **Windows 2000** volume to grow into the adjacent unused space.

| | | | |
|-----------------------------|--|---|--------------------------|
| Disk 0 | | | |
| Basic 40.00 GB Online | DIAGS 39 MB FAT Healthy (EISA Configuration) | Windows 2000 (C:) 30.00 GB NTFS Healthy (System) | 10.00 GB Unused space |

If adjacent space is available, you can immediately use the Virtual Replicator Online Volume Growth feature.

The next diagram shows a configuration with no adjacent, unused space. It is not possible to delete the **DIAGS** volume and grow the **Windows 2000** volume into its place. The only way to free up space for growing **Windows 2000** is by deleting the **OLD TOOLS** volume.

| | | | |
|-----------------------------|---|---|--|
| Disk 0 | | | |
| Basic 40.00 GB Online | DIAGS 39 MB FAT Healthy (EISA Configuration) | Windows 2000 (C:) 30.00 GB NTFS Healthy (System) | OLD TOOLS (D:) 10.00 GB FAT |

When there is no unused adjacent space for a volume to grow into, it may be possible to make space available, depending on your particular system configuration, as explained in the next section.

Adjacent Space on Basic Disks

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

Adjacent Space Through Deletion

If there is another partition already adjacent to the volume that you want 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.

| | | | |
|-----------------------------|---|---|---|
| Disk 0 | | | |
| Basic 40.00 GB Online | DIAGS 39 MB FAT Healthy (EISA Configuration) | Windows 2000 (C:) 30.00 GB NTFS Healthy (System) | 10.00 GB Unallocated [OLD TOOLS (D:)] |

Adjacent Space through Logical Drive Extension

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

| | | |
|-----------------------------|--|---|
| Disk 0 | | |
| Basic 40.00 GB Online | DIAGS 39 MB FAT Healthy (EISA Configuration) | Windows 2000 (C:) 40.00 GB NTFS Healthy (System) |

| | | | |
|--------------------------------|--|---|--------------------------|
| Disk 0 | | | |
| Basic 40.00 GB Online | DIAGS 39 MB FAT Healthy (EISA Configuration) | Windows 2000 (C:) 40.00 GB NTFS Healthy (System) | 40.00 GB Unused space |

| | | |
|--------------------------------|--|---|
| Disk 0 | | |
| Basic 40.00 GB Online | DIAGS 39 MB FAT Healthy (EISA Configuration) | Windows 2000 (C:) 80.00 GB NTFS Healthy (System) |

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 HP 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 | HP ProLiant 1600R |
| Storage Controller | HP Smart Array 3200 |
| External Storage Enclosure | HP 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, for example., database files. |

| Initial Storage Configuration | | | | | | | | | | | | | | |
|-------------------------------|--------|--------|--------|-------|-------|-------|-------|-------|------------------------------|-------|-------|-------|-------|------|
| Channel/I D | 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 | no ne | no ne | non e | | | Bays: 6 internal/14 external | | | | | |
| 1 (External) | 4.3 GB | 4.3 GB | non e | no ne | no ne | non e | no ne | no ne | no ne | no ne | no ne | no ne | no ne | none |

The steps required for performing online volume growth of the example volume are:

- Upgrading the Boot Volume
- Using the ACU to Extend a Logical Drive
- Using Virtual Replicator to Grow the Boot Volume

Upgrading the Boot Volume

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, the upgrade involves replacing the 4.3-GB drives located at ID 0 with 9.1-GB drives.

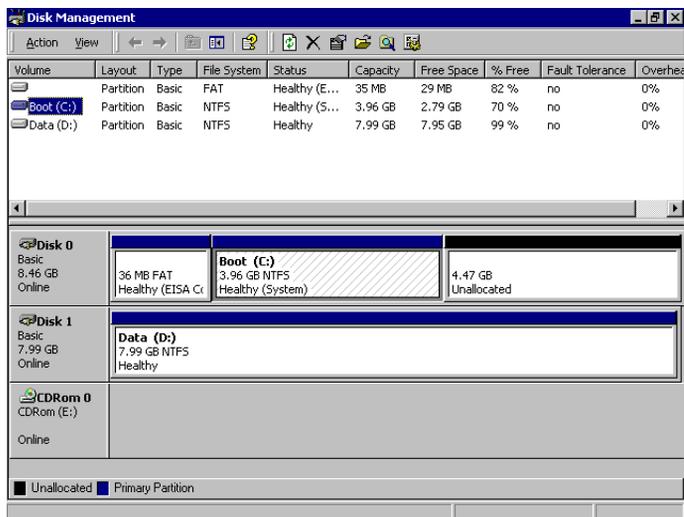
The upgrade results in the following configuration.

| 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 | 9.1 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 |

Using the ACU to Extend a Logical Drive

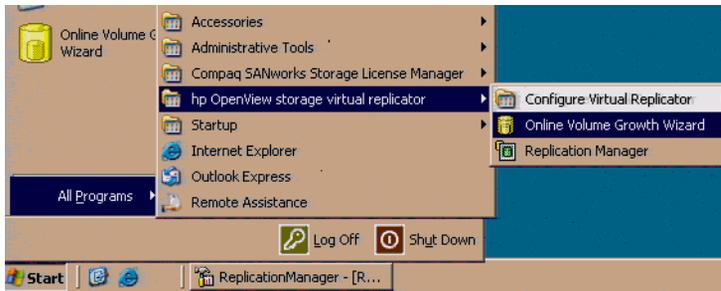
Next, use the HP Array Configuration Utility to perform logical drive extension. By performing this operation, the existing logical drive will grow to use all of the new space provided by the 9.1-GB drives. See the HP Array Configuration Utility documentation for information on extending a logical drive.

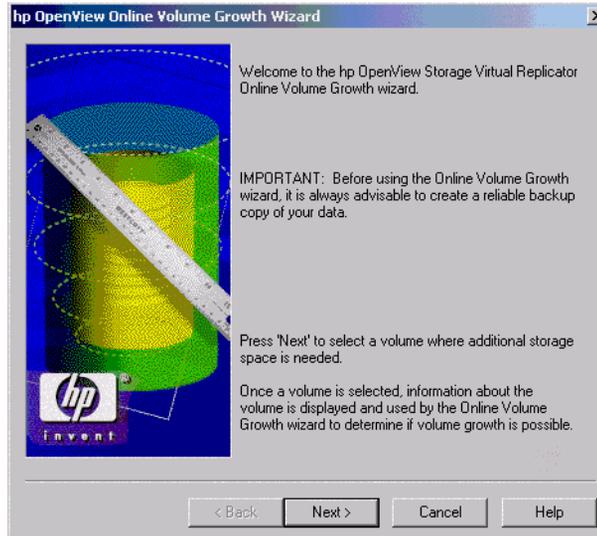
After performing logical drive extension, you can use **Computer Management > Disk Management** to see the basic disk's new layout.



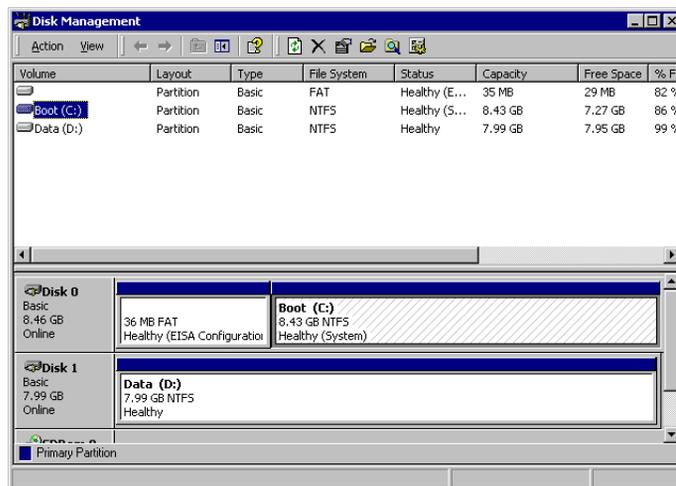
Using Virtual Replicator to Grow the Boot Volume

Now use the **Online Volume Growth Wizard** to create a larger boot volume.





After performing online volume growth, use **Computer Management > Disk Management** to view the resulting layout of the basic disk. The boot volume has grown from 3.96 GB to 8.43 GB without causing any downtime.



Online Volume Growth in Windows Explorer

You can also grow virtual and basic disks through Windows Explorer. The Virtual Replicator Volume Growth wizard is available through 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

See the Virtual Replicator online help for detailed instructions on using Windows Explorer to grow basic and virtual disks.

Scheduling Tasks

10

Virtual Replicator provides a set of wizards that allow you to schedule various virtual disk and snapshot tasks to run automatically. Task scheduling helps you manage recurring operations on virtual disks and snapshots. You have access to the scheduling wizards after you create a virtual disk.

This chapter contains the following topic:

- [Task Scheduling Wizards Overview](#), page 152

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

Task Scheduling Wizards Overview

Virtual Replicator provides the following six wizards for scheduling virtual disk and snapshot tasks:

- **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.

- **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 For Backup**

The Snapshot For Backup wizard schedules a backup operation of a snapshot to tape, including the creation and optional deletion of a snapshot for the backup.

- **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 preset values.

- **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).

Virtual Replicator uses the Windows Scheduled Tasks applet to create and run the schedules for your specified tasks. To use the scheduling wizards, the Scheduled Tasks service needs to be running and you need to have administrative rights.

After scheduling tasks, you can use Virtual Replicator or the Windows Scheduled Tasks applet to view, modify, and delete tasks.

Troubleshooting

11

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

- [General Troubleshooting Questions](#), page 154
- [Incorrect Drive Letters After an Upgrade or Cluster Failover](#), page 155
- [Lost Delayed Write Errors](#), page 156
- [Applications Return Failed Write Errors](#), page 157
- [Accidentally Deleted a Virtual Disk](#), page 158
- [Pool Free Space Fell by More Than the Amount of Data Written](#), page 159
- [Reformatting a Virtual Disk Does Not Work](#), page 160
- [Reconstructing a Pool](#), page 161
- [Pools Not Coming Online](#), page 162
- [Prolonged Startup Times](#), page 163
- [Virtual Replicator Names Are Already in Use](#), page 164
- [Tracking Errors](#), page 165

General Troubleshooting Questions

The following are some general questions you can use when troubleshooting potential VR problems. These questions are also useful when a pool does not come online in a clustered environment.

- What version of VR is running?
- Do you have other resource dependencies defined within the pool group?
- Do you know what caused the original crash on the node? Was it a VR-related crash?
- What type of controllers (such as G80, EVA, and MSA) are you using in your storage?
- Were there any other controller firmware upgrades made to the hardware?
- Are there any errors logged at the controller level?
- When you shut down the node while it was in an online pending state, did it shut down properly?
- When the pool moved back to the peer node, how long did it take to come online?
- When the peer node crashed and the pool was being moved to the other node, were you logged in to the other node?

Note: This last question applies to any 3.x version of Virtual Replicator.

Incorrect Drive Letters After an Upgrade or Cluster Failover

Under some circumstances, virtual disks and snapshots can lose their original drive letter mappings. This can happen when you upgrade Virtual Replicator or when a pool fails over in a cluster. Regarding the pool failover in a cluster, if a drive letter that is mapped to a virtual disk or snapshot is not available on the second node, VR dismounts the drive letter on the second node and reassigns it to the failed over volume. However, if the volume that failed over has the same drive letter as the system disk of the second node, then no drive letter is mapped.

Virtual Replicator provides a Restore Drives utility to help restore mapped drives to virtual disks and snapshots. You can use this utility whenever drive letters become incorrectly mapped. The utility is installed with Virtual Replicator in the default location of `\Program Files\HP\OpenView Storage Virtual Replicator 4.0`.

To run the utility, use a command line, for example:

```
C:\>Program Files\HP\OpenView Storage Virtual Replicator  
4.0\RestoreDrives.exe
```

When you execute the command, the Restore process starts and provides a report about the virtual disks and snapshots that are being remapped. When the process finishes, you will see a Restore Complete message. You can then check for correctly mapped drive letters in the Replication Manager MMC window.

Lost Delayed Write Errors

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 > hp OpenView Storage Virtual Replicator > Replication Manager**.
2. Select the **Replication Manager** snap-in for the computer that is getting the popup errors, so that the snap-in is highlighted.
3. Check the **Free** column in the right-hand Results pane. It shows the free space of all the pools on that computer.

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

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

4. Select the pool that has 0 MB in the **Free** column and expand it to display its virtual disks and snapshots.

Are there any snapshots in the pool?

| | |
|--------|---|
| If Yes | Go to Step 5. |
| If 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.

Note: Deleting files will not increase the pool free space.

- To add a storage unit to the pool, right click the pool, then choose **Properties**. Click the **Storage Units** tab, then click **Add**. Select a storage unit then click **Add**, **OK**, then **OK**. For more information, see [Additional Storage Units in a Pool](#) on page 100.
- To delete a virtual disk or snapshot, right click the virtual disk or snapshot then choose **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 you have applications returning errors when they try to write data to disk, follow these steps to check whether you have a full pool:

1. Click **Start > Programs > hp OpenView Storage Virtual Replicator > Replication Manager**.
2. Select the **Replication Manager** snap-in for the computer that is getting the failed write errors, so that the snap-in is highlighted.
3. Check the **Free** column in the right-hand Results pane. It shows the free space of all the pools on that computer.

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

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

4. Select the pool that has 0 MB in the **Free** column and expand it to display its virtual disks and snapshots.

Are there any snapshots in the pool?

| | |
|--------|---|
| If Yes | Go to Step 5. |
| If 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.

Note: Deleting files does not increase the pool free space.

- To add a storage unit to the pool, right click the pool, then choose **Properties**. Click the **Storage Units** tab, then click **Add**. Select a storage unit then click **Add**, **OK**, then **OK**. See [Additional Storage Units in a Pool](#) on page 100 for more information.
- To delete a snapshot or virtual disk, right click the snapshot or virtual disk, then choose **Delete**. Click **OK** to confirm that you want to delete the disk and erase all the data on it.

Accidentally Deleted a Virtual Disk

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

Note: Deleting a virtual disk will cause the first few blocks of data to be erased. It will be necessary to rebuild these blocks. Contact support for assistance.

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 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.

Note: 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.

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 decrease happens 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, if you write 2 KB of data and 1 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.

Reformatting a Virtual Disk Does Not Work

When you use the Replication 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.

A virtual disk cannot be reformatted if any application is accessing the disk. For example, the problem occurs if you are running Windows Explorer, or if you have a Windows command window open 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 Replication Manager snap-in or command to reformat the disk. For example, close Windows Explorer, or change the prompt in the command window to another drive, then use the snap-in to reformat the disk again.

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.
3. Create virtual disks in the new pool.
4. Use your backup tapes to restore saved data to the new virtual disks.

Pools Not Coming Online

If your pools do not come online, try presenting them one at a time to pinpoint which pools are posing a problem.

If you determine that the problem was caused by a portion of Virtual Replicator being deleted in the registry, you must uninstall Virtual Replicator to clean up the registry. Follow these steps:

1. Shut down one node.
2. From the desktop of the other node, click **Start > Settings > Control Panel > Add/Remove programs**.
3. Select **Virtual Replicator** and remove it from this node.
4. Bring up the node and open **Disk Management**. Verify that the disk drives that make up the pool are available.

Note: If you do not see all of the disks that make up the pool, then Virtual Replicator cannot bind the pool. This means this problem is not related to Virtual Replicator and is really a storage problem.

5. Reinstall Virtual Replicator.
6. Run the namespace recovery utility to get the namespace in the registry restored.
7. If needed, run the `restoreDrives` command.

The existing pools that come back online enable the metadata on the drive to rebuild the pool.

Repeat this procedure on the other node once the procedure has been completed successfully on the first node.

Note: Completing this procedure requires approximately four reboots and takes approximately one hour to finish.

Prolonged Startup Times

If startups or reboots are taking excessive amounts of time, the cause could be that Virtual Replicator pools are requiring longer times to “rebind.” Rebinding occurs during system startup and is the process by which Virtual Replicator pools re-establish the address links between parent disks and snapshots, and load those links into memory.

Normally, rebinding takes little time and has no effect on performance. However, if you have snapshots that have diverged extensively from their parent disks, the time needed to rebind will increase.

Snapshots diverge when changes occur on the parent disk and segments are copied out. Large numbers of copy-outs can occur when you have any of the following:

- Data on parent disks that changes frequently
- Pools with small segment sizes
- Large virtual disks (256 GB to 2 TB) with snapshots present
- Snapshots that are kept for long periods of time.

If the snapshots have many copy-out segments, rebinding will take longer and extend the startup time.

To improve restart times, delete unnecessary snapshots. To prevent this problem in the future:

- Avoid keeping snapshots for long periods of time.
- Use standard backup tools to back up your snapshots.
- Delete snapshots promptly after backing them up.
- When possible, create pools with large segment sizes to reduce copy-out activity.

Virtual Replicator Names Are Already in Use

If you receive a message saying that the name is already in use when you try to create a pool, virtual disk, or snapshot, it might be because the namespace is out of sync.

Virtual Replicator keeps track of the names of pools, virtual disks, snapshots, and logical drives added to a pool in the Windows Registry. If a pool fails, the Replication Manager no longer displays the pool or any of its children. However, the names are still registered in the Windows Registry. If you have namespace errors, use the Namespace Recovery utility with the command:

```
SNAPMGR.EXE UTILITY /RECOVER:NAMESPACE
```

This command will update the registry based on pools that are present and online.

Tracking Errors

Use the Windows 2000/Windows Server 2003 Event Viewer to track errors.

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

- Virtual Replicator Driver events (**cpqvrbus** in Windows 2000/2003)
- Virtual Replicator Lifeguard events (**swvrmon**)

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

Virtual Replicator Driver Events

The **cpqvrbus** (Windows 2000) messages appear in the Event Viewer System Log.

Table 2: Virtual Replicator Driver Event Messages

| Severity | MessageId | Event Source | Event Description |
|----------|-----------|--------------|---|
| Error | 2 | cpqvrbus | The Virtual Replicator driver failed at initialization. |
| Error | 3 | cpqvrbus | Insufficient memory. |
| Error | 4 | cpqvrbus | 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 | cpqvrbus | Virtual Replicator could not read or write to a disk in the pool. |
| Error | 6 | cpqvrbus | 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 | cpqvrbus | Virtual Replicator could not get pointer to device %2, where %2 is the name of the disk device that could not be found. |

| Severity | MessageId | Event Source | Event Description |
|---------------|-----------|--------------|---|
| Error | 9 | cpqvrbus | Free disk space in the pool %2 (where %2 is the name of the pool) is less than 5%. You should delete unneeded snapshots. |
| Warning | 8 | cpqvrbus | Free disk space in the pool %2 (where %2 is the name of the pool) is less than 30%. You should delete unneeded snapshots. |
| Informational | 1 | cpqvrbus | The Virtual Replicator driver has been started. |

Virtual Replicator Lifeguard Events

The **swvrmon** messages are produced by the Virtual Replicator Lifeguard service, which monitors pools. These events are recorded in the Event Viewer Application Log.

Table 3: Lifeguard Event Messages

| 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 | Service handler not installed. |
| 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). |

Table 3: Lifeguard Event Messages

| 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 |

Table 3: Lifeguard Event Messages

| Severity | MessageId | Event Source | Event Description |
|----------|-----------|--------------|--|
| 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 |
| Warning | 4 | swvrmon | <p>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.</p> |

Table 3: Lifeguard Event Messages

| Severity | MessageId | Event Source | Event Description |
|---------------|-----------|--------------|---|
| 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. |

Virtual Replicator Policy Management



By default, Virtual Replicator places certain limits on the creation and use of pools, virtual disks, and snapshots. For example, you can have up to eight virtual disks in a pool and each virtual disk can have up to 12 snapshots. Although you cannot exceed the maximum limits defined by the software, you can fine-tune these and other settings by using Virtual Replicator policies.

The rest of this appendix discusses the following topics:

- [About Policies](#), page 172
- [Policies in the Windows Registry](#), page 173
- [Pool Policies](#), page 175
 - [Maximum Number of Pools \(MaxPools\)](#), page 175
 - [Minimum Free Space for a Pool \(MinFreespacePerPoolPercent\)](#), page 175
 - [Pool Segment Size \(PoolSegmentSize\)](#), page 176
- [Virtual Disk Policies](#), page 177
 - [Maximum Number of Virtual Disks in a Pool \(MaxVirtualPerPool\)](#), page 177
 - [Virtual Disk Allocation Unit Size \(NTFSAllocationUnit\)](#), page 177
- [Snapshot Policies](#), page 179
 - [Maximum Number of Snapshots in a Pool \(MaxSnapshotPerPool\)](#), page 179
 - [Maximum Number of Snapshots of a Virtual Disk \(MaxSnapshotPerFamily\)](#), page 179
- [Policies for the Virtual Replicator Lifeguard](#), page 180

See Appendix B, [Virtual Replicator Lifeguard Service](#), for more information on the Virtual Replicator Lifeguard service.

About Policies

Policies are system settings that help you manage Virtual Replicator resources. Virtual Replicator policies are stored in the Windows Registry and are pre-configured to follow best practices for VR administration. You can modify these pre-set policies to suit your own environment.

You can set policies for pools, virtual disks, and snapshots. All of the policies can be configured to apply to your entire Virtual Replicator environment. Some of the policies can also be set to apply only to individual pools.

For example, if you set a policy that reserves a percentage of a pool's free space for exclusive use by snapshots, you can have the new setting apply to all pools on your system, or only to select pools.

You can use a separate set of policies for the Virtual Replicator Lifeguard service. The Lifeguard policies allow you to monitor storage pools and delete snapshots when pools become full. Lifeguard also monitors pools used for Microsoft Exchange.

On cluster systems, policies are maintained separately for each system and not for the cluster as a whole. Therefore, you need to set the same policies for each node in a cluster.

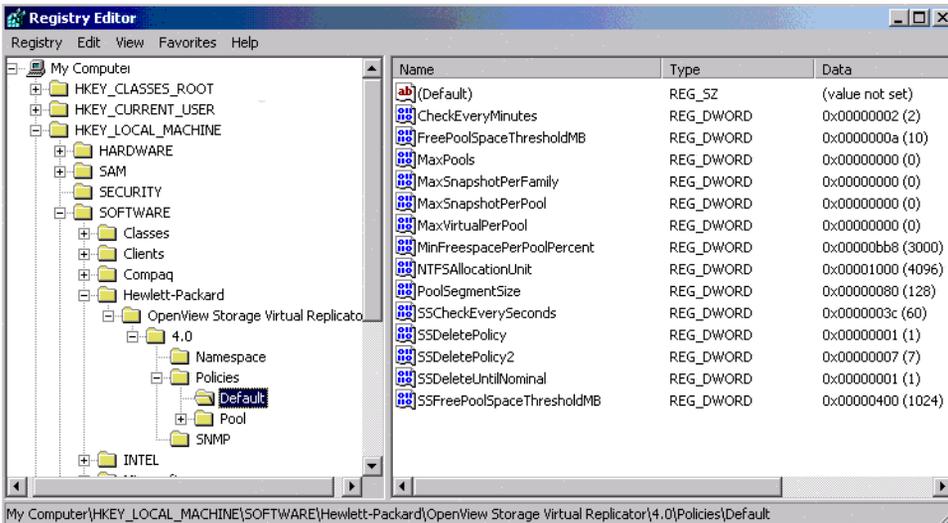
To create or change Virtual Replicator policies, you edit the Windows Registry key values for VR. The instructions for using the Registry Editor to set policies are available in the Virtual Replicator online help.



Caution: Changing VR policy settings requires editing the Windows Registry, which is advised for experienced users only. If there is an error in the Registry, your computer might not function properly. You should be prepared to restore the registry in the event of error. See the Microsoft Registry Editor Help file for more information.

Policies in the Windows Registry

All of the procedures for setting policies begin by starting the Registry Editor and opening the Virtual Replicator registry key. The Registry window looks similar to the next screen:



The right pane of the Registry Editor window shows the following value entries for the Virtual Replicator key:

- CheckEveryMinutes
- FreePoolSpaceThresholdMB
- MaxPools
- MaxSnapshotPerFamily
- MaxSnapshotPerPool
- MaxVirtualPerPool
- MinFreespacePerPoolPercent
- NTFSAllocationUnit
- PoolSegmentSize
- SSCheckEverySeconds
- SSDeletePolicy
- SSDeletePolicy2

- SSDeleteUntilNominal
- SSFreePoolSpaceThresholdMB

You can set policies by double clicking any of the entries in the Registry Editor.

The HKLM...\VR\Policies key contains the default policies that are used by Virtual Replicator. The system-wide policies that apply to all pools are located in the **Policies** subkey.

The HKLM...\VR\Policies\Default key contains the default policies that are only referenced by documentation but are created at installation.

The pre-set default policies are located in the **Policies\Default** subkey.

The policies that apply to specific pools are located in the **Policies\Pool\poolname** subkey. If you need to change a policy for a specific pool, ensure that you select the policy from this path.

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

Pool Policies

Maximum Number of Pools (MaxPools)

MaxPools limits the number of pools that can be created on the system. There is no inherent limit to the number of pools on a system unless this policy is set.

When the number of existing pools equals the policy value, no additional new pools can be created. Note that the count of pools includes “failed” pools.

In a cluster, this policy is not enforced, and there is no limit to the number of pools.

Minimum Free Space for a Pool (MinFreespacePerPoolPercent)

MinFreespacePerPoolPercent maintains a percentage of the pool's free space for exclusive use by snapshots. The default setting is 30%. 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.

If this policy is set to 0, all of the free space in the pool may be used for new virtual disks or for growing existing virtual disks.

The 30% default policy 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. HP strongly recommends that you do not decrease this minimum amount.

You can set this policy for all pools or for individual pools.

The two examples that follow demonstrate how the **MinFreespacePerPoolPercent** policy works.

Example 1:

Pool size = 200 GB

MinFreespacePerPoolPercent = 30 % (60 GB)

Creating a virtual disk of 90 GB would be successful, because it leaves 110 GB of free space. However, creating a second virtual disk of 90 GB would fail because a disk of that size would leave only 20 GB of free space, which is less than 30 percent. You could create a second virtual disk of up to a maximum of 50 GB (assuming there are no snapshots consuming disk space).

Example 2:

Pool size = 150 GB

MinFreespacePerPoolPercent = 30 % (45 GB)

The pool contains a virtual disk of 70 GB and a snapshot consuming 10 GB (leaving 70 GB of free space).

An attempt to grow the virtual disk by 70 GB fails since this would use all the space in the pool. An attempt to grow the virtual disk by 25 GB succeeds because it leaves the minimum of 45 GB of free pool space.

Pool Segment Size (PoolSegmentSize)

Virtual Replicator uses the **PoolSegmentSize** value as the default setting when you create a pool. Segment size determines the largest possible virtual disk or storage unit in a pool.

The pre-set default segment size is 128 KB. This segment size lets you create a virtual disk or add a storage unit up to 1 TB large. The largest segment size available is 256 KB, which allows a virtual disk of 2 TB.

If you change the PoolSegmentSize policy, the new value will be used as the default segment size during pool creation. However, Virtual Replicator will override this policy default if you create a pool and add a large storage unit that requires a larger segment size.

The PoolSegmentSize policy, once set, applies to all pools. You cannot set this policy separately for a single pool.

Virtual Disk Policies

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 eight virtual disks per pool; therefore, only policies of less than eight are useful.

You can apply this policy to all pools on your system or to individual pools.

Virtual Disk Allocation Unit Size (NTFSAllocationUnit)

NTFSAllocationUnit changes the default NTFS allocation unit size for a virtual disk. By default, when formatting virtual disks, Virtual Replicator sets the allocation unit size to 4096 bytes (4 KB). Since, in most cases, this default provides the highest efficiency for storing data, HP recommends that you not change this value.

Note: If you set an allocation unit size greater than 4 KB, you will not be able to defragment the virtual disk.

If you elect to modify this policy, the revised value becomes the new default. This new default value is used when:

- You format a virtual disk using the Replication Manager MMC snap-in and set the Allocation Unit to “Default.”
- You format a disk using the SnapMgr command line interface and do not specify an allocation unit size.

To override the NTFSAllocationUnit policy, you can:

- Pick a nondefault size in the Replication Manager wizard
- Pick a nondefault size through the Command Line Interface
- Use Windows Explorer to format the disk.

Note: This policy has NO effect on other Windows disks—only on virtual disks formatted through Virtual Replicator tools.

Valid values for allocation unit size are (in bytes)

0 (System default is used)

512

1024

2048

4096

8192

16384

32768

65536

You can apply this policy to all pools on your system or to individual pools.

Snapshot Policies

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

MaxSnapshotPerPool sets the maximum number of snapshots that can be created in a pool. Virtual Replicator has an upper limit of 12 snapshots per virtual disk and eight virtual disks per pool. Therefore, only policies of less than 96 are useful.

Note that Virtual Replicator checks the **MaxSnapshotPerPool** and **MaxSnapshotPerFamily** policies simultaneously. Whichever limit is reached first will prevent snapshot creation.

You can apply this policy to all pools on your system or to individual pools.

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

MaxSnapshotPerFamily limits the maximum number of snapshots that can be created within a family. A virtual disk and its child snapshots are known as a family. Virtual Replicator has an upper limit of 12 snapshots per virtual disk (family). Therefore, only policies of less than 12 are useful.

Note that Virtual Replicator checks the **MaxSnapshotPerFamily** and **MaxSnapshotPerPool** policies simultaneously. Whichever limit is reached first will prevent snapshot creation.

You can apply this policy to all pools on your system or to individual pools.

Policies for the Virtual Replicator Lifeguard

Virtual Replicator Lifeguard is a system service that monitors storage pools (see Appendix B, [Virtual Replicator Lifeguard Service](#)). By default, Lifeguard checks for pool-full situations once every 60 seconds, reports on pool conditions in the Windows Application Log, and deletes snapshots when pools become full.

The default snapshot deletion policy is named SSDeletePolicy. This primary policy specifies that when the free space in a pool drops below 1024 MB, Lifeguard begins deleting multiple snapshots, beginning with the oldest, until free space reaches 1024 MB.

If, for any reason, SSDeletePolicy is not carried out, the secondary policy, SSDeletePolicy2, specifies that Lifeguard should delete snapshots with the largest delspace to free up space in the pool.

In a Microsoft Exchange environment, Lifeguard performs a second check of disks every two minutes and shuts down Exchange 5.5 when the storage pool has 10 MB left.

Note: Shutdown of Exchange is not supported on Microsoft Exchange 2000. If you are using Exchange 2000, HP recommends that you set the Lifeguard Snapshot Deletion policies to prevent your pools from running out of free space. See the online help for more information.

You can modify the default Lifeguard frequencies and thresholds by editing the following policies:

- Number of seconds between checks of pool free space (SSCheckEverySeconds)
- Number of minutes between checks of free space in Microsoft Exchange pools (CheckEveryMinutes)
- Threshold at which Lifeguard shuts down Microsoft Exchange (FreePoolSpaceThresholdMB).

You can also modify the primary and secondary snapshot deletion policies by specifying:

- Parameters for the snapshot to be deleted (SSDeletePolicy and SSDeletePolicy2)
- The frequency at which Lifeguard checks the pool (SSCheckEverySeconds)
- The minimum threshold for free space (SSFreePoolSpaceThresholdMB)

- Whether one or multiple snapshots should be deleted at each interval (SSDeleteUntilNominal).

The procedures for modifying primary and secondary snapshot deletion policies are explained in the online help.

Virtual Replicator Lifeguard posts events in the Application Event log, which you can view 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 166 for more information.

You set policies using the Window Registry Editor. See the Virtual Replicator online help for complete instructions on setting policies.

Virtual Replicator Lifeguard Service



Virtual Replicator Lifeguard is a utility that:

- Monitors storage pools and removes snapshots when a pool is nearly full.
- Watches the storage pools used by Microsoft Exchange. On Exchange 5.5, Lifeguard shuts down Exchange if the disks used are in pools that are nearly full.

Note: Automatic shutdown of Exchange is not supported on Microsoft Exchange 2000.

This appendix describes how the Lifeguard service monitors pools.

See Appendix A, [Virtual Replicator Policy Management](#), for information on setting policies for Virtual Replicator Lifeguard.

Monitoring Pool Free Space

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 60 seconds.

By default, when the free space in a pool drops below 1024 MB, Lifeguard begins deleting multiple snapshots, beginning with the oldest, until free space reaches 1024 MB.

If, for any reason, this policy is not carried out, a secondary policy takes effect so that Lifeguard deletes snapshots with the largest delspace to free up space in the pool.

You can change these default policies by modifying the Windows Registry. See [Policies for the Virtual Replicator Lifeguard](#) on page 180 for more information.

Virtual Replicator Lifeguard posts events in the Application Event log, which you can view 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 166 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.

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, Lifeguard checks the amount of current free space in the storage pool that the disk is from. If the amount of free space is less than a specific threshold, Lifeguard gracefully shuts down Microsoft Exchange 5.5. Automatic shutdown is not supported on Exchange 2000.

By default, Virtual Replicator Lifeguard shuts down Microsoft Exchange when the storage pool has only 10 MB left. This check is made once every two minutes and has very little performance impact upon your system. You can change these defaults by creating and/or modifying settings in the system registry. See [Policies for the Virtual Replicator Lifeguard](#) on page 180 for more information.

Virtual Replicator Lifeguard posts events in the Application Event log, which you can view 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. See [Virtual Replicator Lifeguard Events](#) on page 166 for more information.

Checking Pools After Shutdown

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

After you have provided more free space in the storage pool, you can restart Microsoft Exchange.

SnapMgr Commands



This appendix describes the SnapMgr commands that you can use to manage pools, and the virtual disks and snapshots in those pools.

The following table gives a summary of the commands.

| Command | Description |
|-------------|--|
| DRIVES | Shows which drive letters are available to map to virtual disks and snapshots. |
| MANAGE | Controls which computer you manage. |
| 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. |
| UTILITY | Enables support for incremental backups of virtual disks. |
| VIRTUALDISK | Manages virtual disks. |

Use commands at the system prompt or at the **SnapMgr** command prompt. For example, to map the drive letter S to a snapshot called **MySnapshot**, you can type either:

```
C:\> SNAPMGR SNAPSHOT MySnapshot /MAP:S
```

or

```
C:\> SNAPMGR  
SnapMgr> SNAPSHOT MySnapshot /MAP:S
```

If you are at the Windows 2000/Windows Server 2003 command prompt, each command automatically manages the local computer or cluster unless you specify the name of a remote computer after the SNAPMGR command prefix.

For example, this command shows which drive letters are available on the remote computer, Accounts:

```
C:\> SNAPMGR Accounts DRIVES
    Available drive letters: GHKMNPQRST
```

Use the **MANAGE** command to control which computer or cluster you manage.

For example, this command manages the remote computer, Accounts, then shows which drive letters are available:

```
SnapMgr> MANAGE Accounts
SnapMgr> DRIVES
    Available drive letters: GHKMNPQRST
```

For an overview of the command line function, see [Using the Command Line Interface](#) on page 73. The rest of this appendix contains descriptions and examples of all of the SnapMgr commands.

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 and snapshots.

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.

```
SnapMgr> MANAGE HQ
SnapMgr> DRIVES
```

```
Available drive letters: DGHMNPQ
```

2. This example shows which drive letters are available on the standalone computer Backup.

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

```
Available drive letters: DEFGHMNPQRSTVW
```

3. This example shows which drive letters are available on the local cluster node Accounts1.

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

```
Available drive letters: GHKMNPQRST
```

MANAGE

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
```

The MANAGE command is not necessary when you enter commands at the Windows 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 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, such as 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. The node name is required 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 218) to check disk numbers and the POOL command (page 194) to create a pool or add a storage unit to a pool, make sure you specify a node name.

Description

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 prompt will manage the specified standalone computer or cluster.

Examples

This example manages the standalone computer called Accounts. Subsequent commands that you enter at the prompt will manage Accounts.

```
SnapMgr> MANAGE Accounts
```

Manage the Local Computer

Syntax

```
MANAGE /LOCAL
```

Description

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 prompt will manage the local standalone computer or cluster.

Examples

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

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 prompt will manage.

Examples

1. In this example, you are managing the standalone computer called Accounts. Subsequent commands issued at the prompt will manage Accounts.

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

2. In this example, you are managing the Accounts cluster. Subsequent commands issued at the 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
```

- Import a storage unit into a pool:

```
POOL pool /IMPORT: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, up to 23 characters. If you want to use spaces in the name, enclose the name in double 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 218) 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, 256

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.

A pool can have up to eight storage units.

- The 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 storage units should all have the same read-write, redundancy, and failure characteristics. For example, they should all be standard disks, or they should all be RAID 5 arrays, or they should all be mirror sets.

Note: Do not use disks that are part of volume sets, mirror sets, or stripe sets created using Disk Administrator.

- 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 218).



Caution: In a cluster, Virtual Replicator lists *all* storage units as **Shared** in the **Type** column. If you have a logical drive created from available space on a local disk, Virtual Replicator shows it as **Type Shared**, even though it actually is local. It is important that you recognize this type of local storage unit and that you do not add it to a pool; otherwise, the pool will not fail over in the cluster.

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 logical unit that is online and does not contain any partitions. Virtual Replicator does not support removable media. In a cluster, the storage unit 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 is named *poolname* **Group** (where *poolname* is the name you gave the pool). Then it creates a resource of type **SCE Pool** for the pool itself. This resource has the same name as the pool.

The command brings the pool group and pool resource online.



Caution: Do not use Cluster Administrator to rename the pool resource. If you do this, you may not be able to access your data.

You can rename the pool *group*, have more than one pool in the group, and move pools into different groups.

If you create additional resources in the group for applications or file shares that depend on virtual disks or snapshots in a pool, be sure to set cluster dependency on the pool. As a result, the pool will start up before the application or file share starts up.

Examples

1. This example creates a pool called RAID Pool on the HQ cluster. It uses the UNITS command to determine which storage units are available and what their disk numbers are. Because disk numbers can be different on different nodes, the example specifies the name of a node, and not the cluster name, on the MANAGE command. Then it uses the POOL command to create the pool.

```
SnapMgr> MANAGE HQ-NODE1
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, the example specifies the name of a node, and not the cluster name, on the UNITS and POOL commands.

```
C:\> SNAPMGR HQ-NODE1 UNITS
```

| Capacity (MB) | Disk Number | Type |
|---------------|-------------|--------|
| 2007 | 7 | Shared |

```
C:\> SNAPMGR HQ-NODE1 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 218) 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 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 eight storage units in a pool:

- The 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 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 RAID 5 arrays, then add only a RAID 5 array.

Note: Do not use disks that are part of volume sets, mirror sets, or stripe sets created using Disk Administrator.

- 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 218).



Caution: In a cluster, Virtual Replicator lists *all* storage units as **Shared** in the **Type** column. If you have a logical drive created from available space on a local disk, Virtual Replicator shows it as **Type Shared**, even though it actually is local. It is important that you recognize this type of local storage unit and that you do not add it to a pool; otherwise, the pool will not fail over in the cluster.

Each storage unit must be a non-removable logical unit that is online and does not contain any partitions. Virtual Replicator does not support removable media. In a cluster, the storage unit 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.

Examples

1. This example adds a storage unit to the pool named 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, the example 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 HQ-NODE1
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 named 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, the example specifies the name of a node, not the cluster name, on the UNITS and POOL commands.

```
C:\> SNAPMGR HQ-NODE1 UNITS
```

| Capacity (MB) | Disk Number | Type |
|---------------|-------------|--------|
| 4091 | 1 | Shared |

```
C:\> SNAPMGR HQ-NODE1 POOL "News Pool" /ADD:1
```

Import a Storage Unit into a Pool

Syntax

```
POOL pool /IMPORT:disknumber
```

- *pool* is the name of the pool into which you want to import the storage unit.
- *disknumber* is the disk number of the storage unit that you want to import. Use the SNAPINUNITS command to find out which storage units are available and what their disk numbers are.

Description

This command imports a formatted storage unit into an existing pool. Use this command when you want to add disks with pre-existing partitions and data to a pool.

Each logical drive on the disk becomes a new virtual disk in the pool. Any free space on the disk is added to the pool's free space.

The following restrictions apply when you import a storage unit:

- Before importing a storage unit, make sure that the pool has free space equal to or greater than 1% (approximately) of the unit to be imported. This free space is needed for internal configuration data.
- Virtual Replicator does not support the FAT file system. Therefore, if the storage unit you want to import has FAT partitions, you should convert them to NTFS before you import the unit.
- In a cluster, the storage unit to be imported must be a shared disk and must also be a cluster resource.
- On a dual-boot system, the SNAPINUNITS command shows the system disk that **does not contain** the boot partition. Do not import this disk into a pool or you could lose all the data on the disk.

Example

1. This example imports a storage unit into the pool named News Pool. It uses the SNAPINUNITS command to find out which storage units are available and what their disk numbers are. Then it uses the POOL /IMPORT command to import the storage unit.

```
C:\> SNAPMGR SNAPINUNITS
```

```
Capacity (MB)      Disk Number
Type      Partition/s      Free Space
-----
```

```
17360
Num: 1              1          Local
                  16333
```

```
Cap: 1027 MB
Type: NTFS
Drive: F
```

```
C:\> SNAPMGR "News Pool" /IMPORT: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. 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, the command does not show pools that are currently offline.

Summary information

When you omit the *pool* 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: The date 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 222) or the SNAPSHOT command (page 207).

Showing I/O statistics

The /STATISTICS switch shows information about I/Os to the pool. It shows the following counters, which are set to zero 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 number includes reads caused by copy-out operations and split reads. |

| Counter | Description |
|--------------|---|
| Disk writes | The number of write I/Os issued to disk by the pool software. This number includes writes caused by copy-out operations and split writes. |
| Split reads | The number of read I/O requests that had to be split into two or more I/Os to disk because the read crossed a segment boundary, and the next segment was not contiguous with the current segment. |
| Split writes | The number of write I/O requests that had to be split into two or more I/Os to disk because the write crossed a segment boundary, and the next segment was not contiguous with the current segment. |

Showing the storage units

The /UNITS switch shows the following information about each storage unit in a 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 storage unit in a cluster that is on a shared storage bus and can be seen by every node in the cluster.
 - Local:** A physical storage unit 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 | HQ-NODE1 |
| News Pool | 6074 | 3073 | HQ-NODE2 |

- This example shows information about one pool.

```
SnapMgr> POOL "RAID Pool"
```

```
Pool:          RAID Pool
Capacity:      17288 MB
Free space:    8232 MB
Segment size:  32 KB
Created:       11/15/98 08:38 AM
Modified:      1/29/98 08:45 AM
Version:       1.0
Owner node:    HQ-NODE1
```

```
Name          Type      CapacityDel  space  Drive  Family
              (MB)      (MB)
-----
```

```
-----
Accounts Data Virtual  4500  0  Z:1
Engineering Data Virtual 3000 3000Y:2
Accounts MonSnapshot 4500 573W:1
Accounts TueSnapshot4500 2729T:1
```

- This example shows I/O statistics for a pool on the HQ cluster.

```
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
```

- This example shows which storage units are in a pool on the HQ cluster.

```
SnapMgr> MANAGE HQ
SnapMgr> POOL "RAID Pool" /UNITS
```

```
CapacityDisk  Number  Type
      (MB)
-----
```

```
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
```

- Show information about snapshots:

```
SNAPSHOT  
SNAPSHOT snapshot
```

- Assign a mount point to a snapshot:

```
SNAPSHOT snapshot /MOUNT:mount path
```

- Remove a mount point from a snapshot:

```
SNAPSHOT snapshot /UNMOUNT:mount path
```

- List mount points of a snapshot:

```
SNAPSHOT snapshot /LISTMOUNT
```

Create a Snapshot

Syntax

```
SNAPSHOT snapshot /PARENT:parent
```

- *snapshot* is the name that you want to give to the new snapshot, up to 23 characters. If you want to use spaces in the name, enclose it in double 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.

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. The parent disk does not need to 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. The snapshot 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.

Initially, 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.

The new snapshot is initially offline. It remains offline until you map a drive letter to it. Use the SNAPSHOT command (page 207) to map a drive letter to the snapshot, not Disk Administrator or Disk Management.

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. HP recommends this step because during system startup:

- Drive letters are automatically allocated from the beginning of the alphabet.
- The pool software starts late in the startup sequence.

As a result, the drive letter you choose for the snapshot might be allocated to another disk that comes online before the pool software starts. Consequently, the snapshot will have no drive letter assigned. Choosing a letter near the end of the alphabet prevents the allocation of the snapshot's drive letter 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
```

- *snapshot* is the name of the snapshot that you want to delete.

Description

This command deletes a snapshot from a pool.

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 the snapshot's family may increase:

- If the snapshot shared disk space with another snapshot, the Delspace of the other snapshot rises by the amount of the shared disk space.
- If the snapshot 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 does not affect the data stored on its parent disk or on any other disks in its family.

In a cluster, the command fails if the pool cluster resource is offline for the pool that contains the snapshot.

Examples

1. This example deletes the snapshot Accounts Mon on the HQ cluster at the SnapMgr prompt.

```
SnapMgr> MANAGE HQ  
SnapMgr> SNAPSHOT "Accounts Mon" /DELETE
```

2. This example deletes the snapshot Accounts Tue on the HQ cluster at the system prompt.

```
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 contains the snapshot
- Drive letter that is currently mapped to the snapshot
- Family that contains the snapshot
- 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:** The value used by the file system and all other software on the system. For example, Windows Explorer uses the capacity to display the snapshot's size.

Capacity 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 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 | PoolDrive | Family | Owner |
|------------------|-----------|--------|----------|
| Accounts MonRAID | Pool | W: 1 | HQ-NODE1 |
| Accounts TueRAID | Pool | T:1 | HQ-NODE1 |

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:       1/26/01 9:07 PM
Pool:          RAID Pool
```

Assign a Mount Point to a Snapshot

Syntax

```
SNAPSHOT snapshot /MOUNT:mount path
```

- *snapshot* is the name given to the snapshot that needs to be mounted.
- *mount path* is the path to an empty folder on an NTFS volume on which the drive will be mounted.

Description

This command mounts a snapshot onto a specified mount point.

The drive path is persistent; the next time the computer restarts, the same drive path is automatically mapped to the snapshot, provided that the drive path is available.

Note: Assigning mount points to a snapshot is not possible using remote management.

Example

This example mounts a snapshot onto the folder specified.

```
SnapMgr> SNAPSHOT "Sports Snap" /MOUNT:"Z:\Sports Snap"
```

Remove a Mount Point from a Snapshot

Syntax

```
SNAPSHOT snapshot /UNMOUNT:mount path
```

- *snapshot* is the name given to the snapshot that needs to be unmounted.
- *mount path* is the path that needs to be removed.

Description

This command unmounts the selected mount point from the snapshot.

Note: Removing the mount points from a snapshot is not possible using remote management.

Example

This example unmounts the selected mount point from the snapshot.

```
SnapMgr> SNAPSHOT "Sports Snap" /MOUNT:"Z:\Sports Snap"
```

List the Mount Points of a Snapshot

Syntax

```
SNAPSHOT snapshot /LISTMOUNT
```

- *snapshot* is the name given to the snapshot for which all the mount points need to be listed.

Description

This command lists all the mount points of the specified snapshot.

Note: Listing all mount points of a snapshot is not possible using remote management.

Example

This example lists all the mount points available for the snapshot “Sports”.

```
SnapMgr> SNAPSHOT "Sports" /LISTMOUNT
```

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.

The command shows only non-removable logical units 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: a physical disk that is attached to the computer. |

The command also shows disks that are part of volume sets, mirror sets, and stripe sets that were created using Disk Administrator or Disk Management. Do not use these disks in pools.



Caution: In a cluster, Virtual Replicator lists *all* storage units as **Shared** in the **Type** column. If you have a logical drive created from available space on a local disk, Virtual Replicator shows it as **Type Shared**, even though it actually is local. It is important that you recognize this type of local storage unit and that you do not add it to a pool; otherwise, the pool will not fail over in the cluster.

Example

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 HQ-NODE2  
SnapMgr> UNITS
```

| Capacity (MB) | Disk Number | Type |
|------------------|-------------|--------|
| 4091 | 1 | Shared |
| 8678 | 3 | Shared |
| 8678 | 4 | Shared |
| 4091 | 6 | Shared |
| 2007 | 7 | Shared |

UTILITY

Syntax

```
UTILITY /RESET virtualdisk /SNAPSHOT:snapshot  
[/LOGFILE:"logfile_name"]
```

OR

```
UTILITY /RESET virtualdisk /TIME:"MM-DD-YYYY-hh:mm:ss"  
[/MARGIN:mmm] [/LOGFILE:"logfile_name"]
```

- *virtualdisk* is the name of the virtual disk for which you want to turn off archive bits.
- *snapshot* is the name of an existing snapshot that has been backed up.
- *logfile_name* is the name that you want to give to the optional log file.
- *mmm* is the time in minutes for the optional safety margin.

Description

This command enables support for backup tools that perform incremental backups. The Incremental Backup Support utility prevents unnecessary duplicate backups by turning off the archive bit of files on a virtual disk that has a snapshot that has been backed up.

If the snapshot that was backed up still exists, you can use the /SNAPSHOT option, specifying the name of the snapshot.

If the snapshot that was backed up does not exist, use the /TIME option to specify a date and time. The Incremental Backup Support operation changes the archive bits on the files on the virtual disk that are older than or the same as the specified date and time. The parameter format is as follows:

```
"MM-DD-YYYY[-hh[:mm[:ss]]]"
```

Enclose the date/time parameter in double quotes (""). The month, day, and year values are required. Time values are optional and will be set to 0 if omitted. The month, day, hour, minute, and second values can be indicated with 1 or 2 digits. Use a 24-hour format for the hour value. For example, 8-1-2002-23:00:00 indicates August 1, 2002, 11 p.m.

Note: You cannot specify milliseconds as part of the /TIME option. However, Virtual Replicator automatically adds 999 milliseconds to the time you specify. This additional time ensures that files on the virtual disk that were created up to 999 milliseconds after the time you specify have their archive bits turned off.

You can use the /MARGIN option to specify an optional safety margin time in minutes. The number of minutes you specify will be subtracted from the date and time supplied for the /TIME option. The safety margin helps ensure that files that might not have been backed up during an earlier operation are backed up in the next operation. You can set a safety margin of 0 to 999 minutes.

You can use the /LOGFILE option to specify a log file name. The name can be a full path name (with a local drive letter and directory names) or only a file name, and must be enclosed in double quotes (“”). Do not specify a mapped network drive for the log file because the Incremental Backup Support operation may not have access to that drive. If you do not specify a full path name, Virtual Replicator creates the log file in the root directory of your system drive.

Examples

1. This example turns off the archive file bits on the virtual disk Accounts Data, based on the snapshot Accounts Mon, and creates a log file named Accounts Data Log.

```
SnapMgr> UTILITY /RESET "Accounts Data" /SNAPSHOT: "Accounts  
Mon" /LOGFILE:"Accounts Data Log"
```

2. This example turns off the archive file bits of files that are older than August 1, 2002, 11:00 p.m. on the virtual disk mapped to drive letter V:.

```
SnapMgr> UTILITY /RESET V: /TIME: "8-1-2002-23:00:00"
```

VIRTUALDISK

Summary

Use VIRTUALDISK commands to manage virtual disks. You can:

- Assign a mount point to a virtual disk:

```
VIRTUALDISK virtualdisk /MOUNT:mount path
```

- 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
```

- List mount points of a virtual disk:

```
VIRTUALDISK virtualdisk /LISTMOUNT
```

- Map a drive letter to a virtual disk:

```
VIRTUALDISK virtualdisk /MAP:drive
```

- Remove a mount point from a virtual disk:

```
VIRTUALDISK virtualdisk /UNMOUNT:mount path
```

- Restore a virtual disk from an existing snapshot:

```
VIRTUALDISK virtualdisk /RESTORE:snapshot
```

- Show information about virtual disks:

```
VIRTUALDISK  
VIRTUALDISK virtualdisk
```

- Unmap a drive letter from a virtual disk:

```
VIRTUALDISK virtualdisk /UNMAP
```

Assign a Mount Point to a Virtual Disk

Syntax

```
VIRTUALDISK virtualdisk /MOUNT:mount path
```

- *virtual disk* is the name given to the virtual disk that needs to be mounted.
- *mount path* is the path to an empty folder on an NTFS volume on which the drive will be mounted.

Description

This command mounts a virtual disk onto a specified mount point.

The drive path is persistent; the next time the computer restarts, the same drive path is automatically mapped to the virtual disk, provided that the drive path is available.

Note: Assigning mount points to a virtual disk is not possible using remote management.

Example

This example mounts the virtual disk onto the folder specified.

```
SnapMgr> VIRTUALDISK "Sports" /MOUNT:"Z:\Sports"
```

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, up to 23 characters. If you want to use spaces in the name, enclose it in double 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.

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.

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, remaining offline until you map a drive letter to it. Use the VIRTUALDISK command (page 222) 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
```

- *virtualdisk* is the name of the virtual disk that you want to delete.

Description

This command deletes a virtual disk from a pool.

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 there are any snapshots in its family.

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 deletes the virtual disk TempDisk on the HQ cluster at the **SnapMgr** prompt.

```
SnapMgr> MANAGE HQ  
SnapMgr> VIRTUALDISK TempDisk /DELETE
```

2. This example deletes the virtual disk Scratch on the local cluster at the system prompt.

```
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 default allocation unit size is used (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 Virtual Replicator initially sets to 4 KB.

Note: The default setting minimizes the amount of lost space and fragmentation on the volume. HP strongly recommends 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.
- Try to grow a virtual disk over the maximum size of a virtual disk in the pool.
- Try to grow a disk more than the available free space in the pool.

Example

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
```

List Mounts Points of a Virtual Disk

Syntax

```
VIRTUALDISK virtualdisk /LISTMOUNT
```

- *virtualdisk* is the name given to the virtual disk for which all the available mount points need to be listed.

Description

This command lists all the mount points of the specified virtual disk.

Note: Listing all mount points of a virtual disk is not possible using remote management.

Example

This example lists all the mount points available for the virtual disk.

```
SnapMgr> VIRTUALDISK "Sports" /LISTMOUNT
```

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 the virtual disk.

Description

This command maps a drive letter to a virtual disk. The command fails if the pool containing the virtual disk 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. HP recommends this step because during system startup:

- Drive letters are automatically allocated from the beginning of the alphabet.
- The pool software starts late in the startup sequence.

As a result, the drive letter you choose for the virtual disk might be allocated to another disk that comes online before the pool software starts. Consequently, the virtual disk will have no drive letter assigned. Choosing a letter near the end of the alphabet prevents the allocation of the virtual disk's drive letter 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
```

Remove a Mount Point from a Virtual Disk

Syntax

```
VIRTUALDISK virtualdisk /UNMOUNT:mount path
```

- *virtualdisk* is the name given to the virtual disk that needs to be unmounted.
- *mount path* is the path that needs to be removed.

Description

This command unmounts the selected mount point from the virtual disk.

Note: Removing the mount points from a virtual disk is not possible using remote management.

Example

This example removes the selected mount point from a virtual disk.

```
SnapMgr> VIRTUALDISK "Sports" /UNMOUNT:"Z:\Sports"
```

Restore a Virtual Disk

Syntax

```
VIRTUALDISK virtualdisk /RESTORE:snapshot [/POOL:pool]
```

- *virtualdisk* is the name of the virtual disk you want to recover.
- *snapshot* is the snapshot you want to use as the source.
- *pool* is the name of the pool that contains the virtual disk (optional).

Description

This command recreates a virtual disk from the data on a specified snapshot. During the operation, the drive letters of the virtual disk and snapshot are momentarily unavailable.

After the restore operation, the source snapshot is retained, but all previous snapshots of the virtual disk are deleted.

Examples

1. This example restores the virtual disk Accounts Data from the snapshot Accounts Tue on the HQ cluster.

```
SnapMgr> MANAGE HQ  
SnapMgr> VIRTUALDISK "Accounts Data" /RESTORE:"Accounts Tue"
```

2. This example restores the virtual disk Scratch from the snapshot ScratchSnap on the HQ cluster.

```
C:\> SNAPMGR HQ VIRTUALDISK Scratch /RESTORE:ScratchSnap
```

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, the command 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 contains the virtual disk
- Drive letter that is currently mapped to the virtual disk
- Family that contains the virtual disk
- Name of the cluster node that currently owns the 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 | PoolDrive | Family | Owner |
|------------------|---------------|--------|----------|
| Accounts Data | RAID PoolZ:1 | | HQ-NODE1 |
| Engineering Data | RAID PoolY:2 | | HQ-NODE1 |
| Usenet News | News PoolX: 0 | | HQ-NODE2 |

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:     1/26/01 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
```

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.
- *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 you can use to create and 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 of the same virtual disk every day and to keep a week's worth of snapshots. To create one snapshot at a time up to a specified number, you can run SMARTSNAP.EXE manually from the command line. You can also use the **at** command or the Windows Scheduled Tasks applet to schedule SMARTSNAP.EXE to create the specified number of snapshots.

After they are created, SmartSnap snapshots are identifiable by their snapshot name. The name consists of three parts:

- *virtualdisk* - truncated to 17 characters if longer
- SCE~
- a number between 0 and 99

An example snapshot name is: myvdiskSCE~34.

After you set the number of snapshots you want to keep of a given virtual disk, SmartSnap examines all the snapshots and determines if there is room for an additional snapshot.

SmartSnap 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

This example shows how to use the Windows Scheduled Tasks applet to schedule SMARTSNAP.EXE for creating a set of seven snapshots of the virtual disk, myvdisk.

1. Use the Scheduled Tasks applet to add a scheduled task.
2. Browse for **Program Files > HP > OpenView Storage Virtual Replicator > smartsnap.exe**.
3. Set parameters for the task, including name, time, and frequency.
4. After finishing, open the advanced properties dialog box for the task.
5. Add the SmartSnap command parameters after the end of the path name:

```
myvdisk /KEEP:7
```

The task is listed in the **Scheduled Tasks** window. When the task runs at the specified time, it will create one snapshot of myvdisk each day. Before creating the eighth snapshot, SmartSnap will delete the oldest snapshot and then continue.

A

- abbreviating commands [74](#)
- Action menu [67](#)
- adding
 - snap-ins to MMC console files [67](#), [68](#)
 - snapshots to pools [128](#)
 - storage units to pools [100](#), [198](#)
 - virtual disks to pools [114](#)
- adjacent space [143](#) to [145](#)
- Adobe Acrobat Reader [49](#)
- allocation unit [226](#)
- allocation unit size [116](#), [177](#)
- application event log [165](#)
- applications fail or hang [157](#)
- archive bits [63](#), [220](#)
- Array Controllers [144](#)
- assigning mount points
 - to snapshots [215](#)
 - to virtual disks [223](#)
- audience [14](#)
- authorized reseller, HP [19](#)
- automatic snapshots [131](#)
- available drive letters [189](#)

B

- backup
 - command [220](#)
 - incremental [63](#), [220](#)
 - online [29](#)
 - Snapshot for Backup wizard [133](#), [152](#)
 - solutions [133](#)
 - using snapshots for [133](#)

- utilities [63](#)
- basic disks
 - adjacent space [143](#)
 - configuration [138](#)
 - differences with dynamic disks [138](#)
 - partitions [138](#)
 - unsupported volume configurations [139](#)
- benefits of Virtual Replicator [29](#)
- boot drive, upgrading [147](#)
- boot volume, growing [148](#) to [149](#)

C

- cache data [128](#)
- capacity
 - increasing [24](#), [125](#), [137](#)
 - of pools [55](#), [203](#)
 - of snapshots [23](#), [57](#), [208](#), [214](#)
 - of storage units [100](#), [218](#)
 - of virtual disks [22](#), [23](#), [57](#), [114](#), [233](#)
- case insensitivity, in commands [75](#)
- CheckEveryMinutes [180](#)
- Cluster Administrator
 - configuring failback [78](#)
 - restricting pool resource owners [78](#), [99](#)
 - when not to use [76](#), [78](#), [99](#)
- cluster groups
 - creating [78](#), [196](#)
 - deleting [202](#)
- cluster names [62](#)
- cluster resources
 - creating [78](#)
 - pool [99](#)
 - SCE Pool resource [99](#), [196](#), [202](#)

- clusters
 - correcting drive letters after failover 155
 - creating pools 99
 - installing Virtual Replicator 43
 - migrating to 48
 - support for 25
 - using Virtual Replicator on 78
 - commands
 - abbreviating 74
 - DRIVES 189
 - how to use 73, 187
 - MANAGE 190
 - managing remote computers 73
 - POOL 194
 - privileges required 77
 - SMARTSNAP 236
 - SNAPSHOT 207
 - summary of 73
 - UNITS 218
 - upper or lower case in 75
 - UTILITY 220
 - VIRTUALDISK 222
 - Compaq Array Configuration Utility 147
 - console files
 - adding snap-ins to 67, 68
 - default Virtual Replicator files 66
 - explained 66
 - conventions
 - document 15
 - equipment symbols 16
 - text symbols 16
 - copy-out operations
 - and pool rebinding 163
 - and segments 128
 - caused by defragmenting disks 117
 - excessive 163
 - explained 128
 - getting statistics on 110, 204
 - predicting 58
 - cpqvrbus events 103, 165
 - Create Snapshot wizard 152
 - creating
 - cluster groups 196
 - cluster resources 196
 - snapshots 128
 - virtual disks 114
- ## D
- data mining and snapshots 30
 - defragmenting virtual disks 27, 117
 - delayed writes, losing 156
 - Delete Snapshot wizard 152
 - deleting
 - cluster groups 202
 - cluster resources 202
 - pools 202
 - pools in a cluster 112
 - snapshots 212
 - virtual disks 124, 225
 - Delspace
 - explained 105
 - of snapshots 106
 - of virtual disks 105
 - showing 105, 214, 233
 - Device Manager, scanning hardware 76
 - Disk Defragmenter 27
 - Disk Management
 - mounted volumes 26
 - when not to use 76
 - disk space
 - available in a pool 103, 105, 156, 157
 - used by a snapshot 57, 106, 117, 128, 136, 212, 214
 - used by a virtual disk 57, 233
 - disks
 - see snapshots
 - see storage units
 - see virtual disks
 - document
 - conventions 15
 - related documentation 14
 - documentation 49
 - drive letters
 - correcting 45, 155

- mapping 115
 - to snapshots 210
 - to virtual disks 113
- showing available 189
- unmapping 115
 - from snapshots 211
 - from virtual disks 234
- DRIVES command 189
- dynamic disks
 - configuration 140
 - differences with basic disks 140
 - extending volumes 140
 - fault tolerance 140
 - logical drive configuration 140
- E**
- equipment symbols 16
- error messages
 - cpqvrbus events 165
 - failed write 157
 - lost delayed write 156
 - pool 70% full 103
 - pool 95% full 103
 - sddriver events 165
 - swvrmon events 166
- event log 103, 165
- F**
- family
 - explained 128, 224
 - showing 118, 134, 213, 232
- FAT file system
 - converting 101
 - not supported 56, 141
- fault tolerance
 - dynamic disks 140
 - supported 100
- formatting virtual disks 116, 160, 226
- free space in pools
 - about 103
 - and Delspace 105
 - increasing 136
 - monitoring 103
 - specifying 172, 175
- FreePoolSpaceThresholdMB 180
- G**
- getting help 19
- growing
 - basic disks 143 to 149
 - virtual disks 142, 227
- H**
- help, obtaining 19
- HP
 - authorized reseller 19
 - storage website 19
 - technical support 19
- HP Array Configuration Utility 144
- I**
- I/O statistics, for pools 110, 204
- import unit 100, 200
- Incremental Backup Support 63, 220
- installation
 - components 36
 - in a cluster 43
 - Internet Explorer requirements 34
 - operating system requirements 34
 - prerequisites 34
 - procedure 38
 - Snapshot Planner 37
- Internet Explorer, installation requirement 34
- L**
- licenses
 - install FLEXlm on dedicated server 42
 - retrieving key 35
 - setting up 40
 - Storage License Manager 40
- Lifeguard service 183 to 185
- Microsoft Exchange pools 180, 185

- policies [180](#)
- listing mount points
 - of snapshots [217](#)
 - of virtual disks [228](#)
- log, error [103](#)
- Logical Disk Manager [76](#)
- logical drives
 - dynamic disks [140](#)
 - extending [144](#), [147](#)

M

- MANAGE command [190](#)
- Management Information Base [25](#)
- mapping drive letters
 - to snapshots [210](#)
 - to virtual disks [113](#), [229](#)
- MaxPools [175](#)
- MaxSnapshotPerFamily [179](#)
- MaxSnapshotPerPool [179](#)
- MaxVirtualPerPool [177](#)
- MIB [25](#)
- Microsoft applications [27](#)
- Microsoft Exchange [27](#)
 - automatic shutdown [180](#), [185](#)
 - monitoring pool free space [180](#), [185](#)
- Microsoft Management Console
 - see MMC
- MinFreespacePerPoolPercent [175](#)
- mirrorsets
 - restrictions [53](#), [56](#), [139](#)
 - storage units [100](#)
 - using as storage units [53](#)
- MMC
 - Action menu [67](#)
 - adding snap-ins [68](#)
 - console files [66](#)
 - Results pane [67](#)
 - Scope pane [66](#)
 - shortcut menus [67](#)
 - understanding [66](#)
 - View menu [67](#)
- monitoring

- I/Os to pools [110](#), [204](#)
- pool free space [103](#), [184](#)
- monitoring performance [26](#)
- mount points
 - assigning
 - to snapshots [215](#)
 - to virtual disks [223](#)
 - listing
 - of snapshots [217](#)
 - of virtual disks [228](#)
 - removing
 - of snapshots [216](#)
 - of virtual disks [230](#)
- mounted volume [26](#)
- .msc files [66](#)
- multiple partitions, not supported [56](#)
- My Computer [119](#)

N

- names
 - cluster names [62](#)
 - uniqueness [62](#)
- namespace recovery [164](#)
- NTFSAllocationUnit [177](#)

O

- online backup [29](#), [133](#)
- online volume growth [24](#)
 - basic disks [138](#)
 - boot volume, growing [148](#) to [149](#)
 - creating adjacent space [144](#)
 - dynamic disks [140](#)
 - example [145](#) to [149](#)
 - HP Array Configuration Utility [144](#)
 - partitions [139](#)
 - permissions needed [141](#)
 - preparing for [143](#)
 - storage units in pools [142](#)
 - unsupported configurations [139](#)
 - using Windows Explorer [150](#)
 - virtual disks [142](#)

Oracle 27

P

partitioning virtual disks 116

partitions

backwards compatibility 139

basic disks 138

deleting to create space 144

FAT, not supported 56

importing storage units 100

multiple, not supported 56

online volume growth 139

Performance Monitor 26

policies

about 172

CheckEveryMinutes 180

deleting snapshots 180, 184

FreePoolSpaceThresholdMB 180

MaxPools 175

MaxSnapshotPerFamily 179

MaxSnapshotPerPool 179

MaxVirtualPerPool 177

NTFSAllocationUnit 177

pools 175

PoolSegmentSize 176

preventing full pools 180

snapshots 179

SSCheckEverySeconds 180

SSDeletePolicy 180

SSDeletePolicy2 180

SSDeleteUntilNominal 180

SSFreePoolSpaceThresholdMB 180

virtual disks 177

POOL command 194

pools

adding partitioned storage units to 100

adding storage units to 100, 198

benefits of 29

capacity of 55, 203

cluster resources 78, 99

creating in a cluster 99

deleting 202

deleting in a cluster 112

failure 53

free space 136, 172, 175

getting full 103, 156, 157, 180

importing storage units 100, 200

maximum number of 175

maximum number of snapshots 179

maximum number of virtual disks 114, 177

maximum size of 55

Microsoft Exchange pools 185

monitoring free space 103, 184

names of 62

prerequisites 100

properties 108

rebinding 163

reconstructing 161

removing storage units from 102

requirements 100

segment size 176

setting policies for 175, 180

showing information about 108, 203

specifying free space 172, 175

specifying segment size 176

storage unit requirements 53

supported disks 100

troubleshooting free space 159

PoolSegmentSize 176

possible owners

in clusters 78

privileges 77, 141

properties

of snapshots 134

of virtual disks 118

R

rack stability, warning 18

RAID

software RAID, not supported 56

rebinding of pools 163

reconstructing a pool 161

Registry Editor 173

related documentation 14

- remove unit [102](#)
- removing
 - snapshots from pools [212](#)
 - virtual disks from pools [225](#)
- removing mount points
 - of snapshots [216](#)
 - of virtual disks [230](#)
- Restore Drives utility [45](#), [155](#)
- Restore from Snapshot wizard [152](#)
- restoring
 - drive letters [45](#), [155](#)
 - Restore from Snapshot wizard [152](#)
 - solutions [133](#)
 - virtual disks [121](#) to [123](#)
- restrictions
 - mirrorsets [53](#), [56](#), [139](#)
 - on use of virtual disks and snapshots [56](#)
 - storage units [53](#)
 - stripesets [53](#), [56](#), [139](#)
 - volume sets [53](#), [56](#), [139](#)
- Results pane of MMC [67](#)
- right-click menus [67](#)
- S**
- SCE Pool resources
 - creating [196](#)
 - deleting [112](#), [202](#)
- scheduled tasks
 - deleting [152](#)
 - modifying [152](#)
 - viewing [152](#)
- scheduling
 - automatic snapshots [152](#)
- Scope pane of MMC [66](#)
- sddriver events [103](#), [165](#)
- security
 - privileges for online volume growth [77](#), [141](#)
 - privileges for snap-ins and commands [77](#)
 - SWVR Admins group [77](#)
- segment size [176](#)
 - policy [176](#)
 - setting [83](#)
- segments
 - and copy-outs [128](#)
 - and pool free space [159](#)
 - and pool rebinding [163](#)
 - copying out [159](#)
 - explained [128](#)
- shortcut menus [67](#)
- shortening commands [74](#)
- showing
 - available drive letters [189](#)
 - available storage units [218](#)
 - information about pools [108](#), [203](#)
 - information about snapshots [213](#)
 - information about virtual disks [232](#)
- size
 - see capacity
- SMARTSNAP utility [236](#)
- SnapBack [121](#) to [123](#)
- snap-ins
 - adding [67](#), [68](#)
 - explained [66](#)
 - privileges required [77](#)
 - using [65](#)
- SnapMgr commands
 - abbreviating [74](#)
 - DRIVES [189](#)
 - how to use [73](#), [187](#)
 - MANAGE [190](#)
 - managing remote computers [73](#)
 - POOL [194](#)
 - privileges required [77](#)
 - SNAPSHOT [207](#)
 - summary of [187](#)
 - UNITS [218](#)
 - upper or lower case [75](#)
 - UTILITY [220](#)
 - VIRTUALDISK [222](#)
- SNAPSHOT command [207](#)
- Snapshot for Backup wizard [133](#), [152](#)
- Snapshot Planner
 - installation [37](#)
 - uninstalling [58](#)

- using 58
 - Snapshot Watchdog wizard 152
 - snapshots
 - assigning mount points to 215
 - benefits of 29
 - best practices 130
 - capacity 23
 - creating 128
 - creating automatically 131
 - deleting 136, 212
 - deletion policies 180, 184
 - delspace 106
 - disk space used by 57, 58, 106, 117, 128, 136, 212, 214
 - explained 128 to 129
 - in a cluster 130
 - keeping 131
 - listing mount points of 217
 - mapping drive letters to 115, 210
 - maximum number in a pool 179
 - names of 62
 - number of 128, 179
 - planning 57, 58
 - removing mount points of 216
 - restrictions 53, 56
 - scheduling tasks 152
 - setting policies for 179
 - showing information about 213
 - unmapping drive letters from 115, 211
 - using for backup 133
 - using to restore virtual disks 121 to 123, 133
 - SNMP
 - agent 25
 - support for 25
 - SQL 27
 - SSCheckEverySeconds 180
 - SSDeletePolicy 180
 - SSDeletePolicy2 180
 - SSDeleteUntilNominal 180
 - SSFreePoolSpaceThresholdMB 180
 - startup times 163
 - statistics, I/Os to pools 110, 204
 - storage capacity
 - increasing 24
 - Storage License Manager
 - creating license files 40
 - other tasks 41
 - overview 40
 - storage pools
 - see pools
 - storage units
 - adding to pools 100, 198
 - characteristics 100
 - importing 100, 200
 - in clusters 100
 - removing from pools 102
 - requirements for inclusion in a pool 100
 - shared 196, 199
 - showing all available 218
 - showing which are in a pool 109, 205
 - supported disks 100
 - with partitions 100
 - stripesets
 - restrictions 56, 139
 - using as storage units 53
 - SWVR Admins group 77
 - swvrmon events 165, 166
 - symbols in text 16
 - symbols on equipment 16
 - system event log
 - cpqvrbus events 165
 - pool 70% full 103
 - pool 95% full 103
 - sddriver events 165
- ## T
- task scheduling 151 to ??
 - Task Scheduling Wizards 152
 - technical support, HP 19
 - testing applications 30
 - text symbols 16
 - troubleshooting 153

U

UNITS command [218](#)
unmapping drive letters
 from snapshots [211](#)
 from virtual disks [234](#)
unmounting paths
 to snapshots [216](#)
 to virtual disks [230](#)
upgrading
 and migrating to Windows 2000 [46](#)
 correcting drive letters [45](#), [155](#)
 from evaluation copy [47](#)
 from Version 2.0 [44](#)
user interfaces [28](#)
UTILITY command [220](#)

V

View menu [67](#)
virtual disks
 allocation unit size [177](#)
 assigning mount points to [223](#)
 benefits of [29](#)
 capacity [22](#), [23](#), [114](#)
 creating [114](#)
 defragmenting [117](#)
 deleting [124](#), [225](#)
 delspace [105](#)
 disk space used by [57](#), [233](#)
 features [22](#)
 growing [142](#), [227](#)
 listing mount points of [228](#)
 mapping drive letters to [113](#), [115](#), [229](#)
 maximum number in a pool [52](#), [114](#), [177](#)
 maximum number of snapshots [52](#), [128](#), [179](#)
 maximum size of [52](#)
 names of [62](#)
 online volume growth [142](#)
 partitioning [160](#)
 partitioning and formatting [116](#), [226](#)
 planning [57](#)
 policies [177](#)

 properties [118](#)
 reformatting [160](#)
 removing mount points of [230](#)
 restoring [121](#) to [123](#)
 restrictions [53](#), [56](#)
 scheduling tasks [152](#)
 setting policies for [177](#)
 showing information about [118](#), [232](#)
 unmapping drive letters from [115](#), [234](#)
 using as mount points [26](#)
 volume label [116](#)
Virtual Replicator
 using with backup tools [63](#)
 using with other tools [76](#)
VIRTUALDISK command [222](#)
volume label [93](#), [116](#)
volume mount points [26](#)
volume sets
 restrictions [56](#), [139](#)
 using as storage units [53](#)
volumes
 creating adjacent space [144](#)
 growing [24](#), [143](#), [148](#)

W

warning
 rack stability [18](#)
 symbols on equipment [16](#)
websites
 HP storage [19](#)
Windows 2000
 installation requirement [34](#)
 migration from NT 4.0 [46](#)
Windows Explorer
 growing disks [150](#)
 restriction [76](#), [104](#)
 viewing disk properties [104](#)
 volume labels [93](#)
Windows NT, installation requirement [34](#)
Workday Snapshot wizard [152](#)