System Administrator's Guide

HP OpenView Storage Virtual Replicator

Product Version: 4.1

Fourth Edition (July 2004)

Part Number: AA–RENKH–TE

This guide describes how to use Storage Virtual Replicator (SVR) for storage virtualization, online volume growth, and snapshots.

Storage Virtual Replicator System Administrator’s Guide
Fourth Edition (July 2004)
Part Number: AA–RENKH–TE

contents

About This Guide

This guide provides information to help you:

- Plan your Storage Virtual Replicator (SVR) implementation.
- Understand the SVR storage management features.
- Get started using the SVR software.

“About This Guide” describes the following topics:

Overview

This section covers the following topics:

Intended audience

Readers must be familiar with SVR as well as the following:

- Microsoft Windows 2000 and 2003 servers and workstations
- Microsoft Cluster Server (MSCS) clusters

Related documentation

Refer to the following documentation for more information about SVR:

- *HP OpenView Storage Virtual Replicator Command Line Interface Reference Guide*
- *HP OpenView Storage Virtual Replicator Installation Instructions*
- *HP OpenView Storage Virtual Replicator Release Notes*
- HP OpenView Storage Virtual Replicator Online Help
- HP OpenView Storage Virtual Replicator Online Volume Growth Online Help
- HP OpenView Storage Virtual Replicator Scheduling Wizards Online Help
- HP OpenView Storage Virtual Replicator Planning Charts

Conventions

Conventions consist of the following:

- Document conventions
- Text symbols
- Equipment symbols

Document conventions

This document follows the conventions in Table 1.

Table 1: Document conventions

Convention	Element
Blue text: Figure 1	Cross-reference links
Bold	Menu items, buttons, and key, tab, and box names
Italics	Text emphasis and document titles in body text
`Monospace font`	User input, commands, code, file and directory names, and system responses (output and messages)
`Monospace, italic font`	Command-line and code variables
Blue underlined sans serif font text (http://www.hp.com)	Web site addresses

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.

Tip: Text in a tip provides additional help to readers by providing nonessential or optional techniques, procedures, or shortcuts.

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 injury 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 injury 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 injury 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 injury 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 injury 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 web site: http://www.hp.com.

HP technical support

Telephone numbers for worldwide technical support are listed on the following HP web site: http://www.hp.com/support/. From this web site, 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 web site

The HP web site 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 web site, 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 web site for locations and telephone numbers: http://www.hp.com.

Storage Virtual Replicator Overview

1

This chapter provides an overview of Storage Virtual Replicator (SVR). It describes the following topics:

Introduction

Storage Virtual Replicator (SVR) enables you to centrally manage storage virtualization in Windows environments. You can simplify storage configuration and management and enhance availability and scalability.

SVR 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 virtual disks.
- Use the virtual disks and snapshots on your local computer.

The SVR software runs on stand–alone computers and on computers in Microsoft Cluster Server (MSCS) clusters.

Note: Unless otherwise noted, all future references to Windows 2000 and Windows Server 2003 will be identified collectively as Windows.

Benefits

Using SVR, you can:

- Optimize disk storage through storage pooling and virtualization.
- Add storage capacity and expand volumes without rebooting.
- Use virtual disks to allocate disk space easily.
- Create snapshots to:
  - Back up your data online.
  - Maintain copies of your data online and perform quick file restores.
  - Test applications.
- Manage your environment remotely.

Installation

For instructions on installing SVR, refer to the *HP OpenView Storage Virtual Replicator Installation Instructions*.

Features

SVR has several features that enable you to configure and manage your storage environment. This section describes the following topics:

- Storage pooling
- Virtual disks
- Snapshots
- Snapshot task scheduler
- Data migration
- Online Volume Growth
- Interfaces

Storage pooling

You can group hardware array storage, or physical disks, into a logically concatenated *pool* of disk space. You can create an unlimited number of pools using any storage to which Windows systems have direct access. SVR supports both single-spindle disks and controller-based, fault-tolerant disk arrays, such as HP StorageWorks RAID arrays, often called *storage units*.

See “Managing Pools” on page 41 for more information.

Virtual disks

In a pool, you can create virtual disks to tailor disk space to the needs of your users and their applications. For example, a user needs 650 MB of disk space so you create a 650 MB virtual disk. To create 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 the physical storage. Virtual disk size ranges from 10 MB to 2 TB and depends on the pool’s free space, segment size, and policies.

Virtual disks function as physical disks. You can format a virtual disk, map drives to it, read and write data to/from a virtual disk, and install applications on it.

See “Managing Virtual Disks” on page 59 for more information.

Snapshots

You can create instant copies, called *snapshots*, of virtual disks and use them as you would physical disks. You can read from snapshots and write to them, and you can create snapshots while users are reading from and writing to the original virtual disk.

See “Managing Snapshots” on page 81 for more information.

Snapshot task scheduler

You can use the SVR scheduling wizards to automatically create, delete, back up, and restore snapshots. You can specify the frequency and time that each task is performed.

See “Scheduling Tasks” on page 119 for more information.

Data migration

You can migrate data from one physical disk to another without interrupting service. This may be useful when:

- Migrating from one storage system to another during work hours
- Replacing obsolete, slow storage with higher capacity, quicker-spinning storage
- Adding higher capacity storage to pools that have exceeded their limits

See “Migrating Data” on page 99 for more information.

Online Volume Growth

You can use the SVR Online Volume Growth (OVG) wizard to increase storage capacity without disrupting operations on Windows systems. Typically, when you expand a RAID set, the operating system does not recognize the size change until you reboot. However, OVG directs the operating system to update the size of a physical or virtual disk without rebooting.

You can also use OVG to grow the on-disk partition information for a volume. You can expand a volume into unused space made available by either increasing the size of a disk or by deleting one or more adjacent partitions.

See “Online Volume Growth” on page 105 for more information.

Interfaces

You can access SVR tools using the Microsoft Management Console interface (MMC) or the command line interface (CLI). Use either interface to monitor and manage your storage resources.

See “Getting Started with Storage Virtual Replicator” on page 29 for more information.

Supported tools

This section describes the tools you can use with SVR:

- Cluster support
- SNMP support
- Windows tools support
- Interoperability with other storage management tools

Cluster support

SVR works with 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.

You can use the SVR management tools to create cluster resources and simplify cluster management. See “Using cluster resources” on page 40 for more information.

SNMP support

You can use SVR with the Simple Network Management Protocol (SNMP) to exchange management information between management console applications (such as UniCenter) and managed entities. When you install SVR, you also install:

- **SNMP agent**—A processing element that retrieves local management information based on requests from applications monitoring the nodes.
- **Management Information Base (MIB)**—A collection of managed objects in a database that defines the variables in the tables, the data to be retrieved, and the format in which data is presented.

The SNMP agent ensures data is returned to the node that requested it. The MIB is installed when you install SVR on the local computer. The default location is:

```
\Program File\HP\OpenView Virtual Replicator\*.mib
```

Locate the MIB file and copy it to the node on which your management console applications reside.

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

Windows tools support

You can use SVR with the following Windows tools:

- Windows Performance
- Disk Management
- Disk Defragmenter

Windows Performance

Use the Windows Performance tool to monitor virtual disks and snapshots. The tool includes disk performance counters, such as Disk Read Bytes/sec and Disk Write Bytes/sec and are listed under the **PhysicalDisk** or **LogicalDisk** performance object.

SVR also adds the SWVR_Pool, SWVR_VirtualDisk, and SWVR_Snapshot objects so you can monitor SVR performance activity. For example, you can monitor the available free space in a pool or the Delspace of a snapshot.

Disk Management

You can use the virtual disks that you create with SVR as Windows volume mount points. Use the Disk Management tool to establish 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 26 drive letters.

Note: Windows 2000 cluster systems do not support volume mount points.

Several SVR features, such as Incremental Backup Support, Online Volume Growth, and virtual disk formatting, require that virtual disks have a mapped drive letter.

For more information about mounted volumes, refer to your Windows documentation.

Disk Defragmenter

Use the Disk Defragmenter tool to defragment virtual disks and increase I/O performance. This tool finds and consolidates fragmented files and folders on a volume. It also consolidates free space, making it less likely that new files will be fragmented. As a result, your system can access files and folders and save new ones more efficiently. See “Defragmenting virtual disks” on page 70 for more information.

Interoperability with other storage management tools

SVR works with other HP storage applications, such as HP StorageWorks Secure Path and HP StorageWorks Command Console (SWCC), as well as storage applications from other vendors. See the *HP OpenView Storage Virtual Replicator Release Notes* for information about other storage applications.

You can also use SVR with backup tools from other vendors, such as Computer Associates ArcServe and VERITAS Backup Exec, as well as database and messaging applications, such as Microsoft SQL Server and Microsoft Exchange.

Planning Storage Growth

2

This chapter describes how to plan storage for a stand–alone configuration and a cluster configuration:

Planning overview

When planning storage growth for your configuration, consider these SVR guidelines:

- **Pools**—The number of pools is limited only by the number of available storage units. (See "Using storage units in pools" on page 23 for more information.)
- **Virtual disks**—Each pool can have a maximum of 8 virtual disks. The capacity of each virtual disk ranges from 10 MB to 2 TB. The disk's capacity is limited by the pool's free space, segment size, and policies.
- **Snapshots**—Each *family* can have a maximum of 12 snapshots. A family comprises a virtual disk, its snapshots, and snapshots of those snapshots. For example, a family can have 12 snapshots of each virtual disk, or it can have 11 snapshots of a virtual disk, with one snapshot having its own snapshot.

Use the planning charts included on the Virtual Replicator installation CD. One is an example of a completed chart (`Example planning chart.pdf`) and one is blank for you to complete (`Planning chart.pdf`).

Pools

This section describes planning guidelines and recommendations for pools:

Using storage units in pools

HP recommends that all storage units in a pool have the same redundancy, read-write, and failure characteristics to ensure:

- **Consistent results**—Mixing different types of storage units in the same pool may cause unpredictable results. For example, storage units in a pool should be all HP StorageWorks RAID 5 storage sets, all mirror sets, or all standard disks.
- **Single point of failure**—A pool is a single point of failure. If one of the pool's storage units becomes inaccessible, you lose the pool and all of its virtual disks and snapshots. For example, you should not use a standard single-spindle disk and a mirror set (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 pool if the standard disk fails.

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.

If you want the data on a virtual disk to have particular characteristics, create it in a pool whose storage units have those characteristics. For example, if you want the data on a virtual disk to be mirrored, create the virtual disk in a pool whose storage units are all controller-based mirror sets.

If you lose a pool, you have to reconstruct it. To do this, create a new pool, create new virtual disks in the pool, and then restore data from your backup tapes to the new virtual disks.

Storage unit guidelines

When deciding which storage units to use in a pool:

- Create pools from whole, unformatted disks, not from partitions or logical drives on a disk.
- A pool can contain up to 12 storage units. They can have different capacities and be from different manufacturers, but they must be the same type.
- In a cluster, use disks on the shared storage bus only; 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.

Disk space and pools

When a pool runs out of disk space, writes to the pool fail. How soon you are notified of the failed writes depends on whether the applications using the pool's disks use the system cache. Applications that do not use the system cache report failed writes immediately.

Most applications use the system cache and do not report that the pool is full until the system cache tries to write data to disk. A message is displayed indicating that a delayed write failed and data may be lost.

HP recommends that you plan accordingly so pools do not run out of disk space. Disk space becomes an issue if:

- A pool contains snapshots. Be sure to monitor the pool's free space regularly. (See "Monitoring pool free space" on page 49 for more information.)
- A pool's capacity is less than the total capacity of its virtual disks and snapshots.
- You keep snapshots longer than necessary, which increases the time needed to restart the system.

To increase a pool's disk space:

- Add a storage unit to the pool. You can do this dynamically while users are reading from and writing to the pool's virtual disks and snapshots.
- Delete one of the pool's snapshots or virtual disks.

Caution: Do not try to increase a pool's disk space by deleting files. Deleting files uses more space by copying segments of data to snapshots.

Pool capacity

A pool can contain up to 12 logical storage units. Your hardware and software configurations determine the maximum size of the pools you create. For example, you have an HP HSG80 controller with ACS version 8.6 firmware. You can configure storage units of 2 TB each and create a maximum pool size of 24 TB. Different configurations yield different results.

SVR supports pools with a maximum segment size of 256 KB. You create a virtual disk with a maximum size of 2 TB if the segment size is 256 KB.

The actual pool capacity is less than the combined capacity of its storage units because the pool configuration data uses disk space. The amount of space that the configuration data uses will vary and can be up to 10 percent of the combined capacity of the storage units.

Planning the pool size

You can create a large pool that contains several virtual disks, or you can create several small pools, each of which contains fewer virtual disks.

When planning the pool size, consider the following:

- Determine the applications you want in each pool. You can assign each application its own virtual disk, rather than have applications share virtual disks.
- A pool is a unit of failover in a cluster. If the node on which a pool resides fails, the pool and its virtual disks fail over to another node in the cluster. All applications must fail over to the pool at the same time.

Using large pools

Use large pools with several virtual disks to:

- **Protect against hardware failures.** Use storage units that provide redundancy, such as HP StorageWorks RAID 5 storage sets.
- **Reduce the cost of disk space for snapshots.** For example, you have three 10 GB virtual disks and only need snapshots to back up the disks. The snapshot requires 10 GB because it must be at least the same capacity as the virtual disk.

  If you want to back up all disks simultaneously, create one 40 GB pool. Within the pool, create three 10 GB virtual disks, which leaves 10 GB of free space for a snapshot (the virtual disks "time-share" the free space for their snapshots). This saves 50% more disk space than if you created a pool for each disk, which would require three 20 GB pools.
- **Manage several applications that require minimal disk space.** and you have a large disk, create a pool from the large disk and divide it into several small virtual disks.
- **Save disk space if you are not using snapshots.** For example, you have five applications, each of which requires 10 GB of disk space, but you have only 20 GB disks. You can combine three of the disks to make one pool. The capacity of the pool will be approximately 54 GB (about 10 percent less than the combined capacity of the storage units), which is enough for five 10 GB virtual disks.

Using small pools

Use small pools to:

- **Balance the load manually and control which applications run on the nodes.** If you have one pool for each application, you have maximum control.
- **Balance the load across individual storage units.** However, if a pool contains multiple units, SVR uses a best–fit algorithm to determine which LUN to use for mapping the space required by the virtual disk.

Virtual disks and snapshots

This section describes guidelines for using virtual disks and snapshots:

Exceptions

You can use a virtual disk or snapshot as you would a physical disk except in the following situations:

- You cannot use a virtual disk or snapshot as your system disk.
- You cannot create more than one partition on a snapshot. However, you partition a large disk into smaller disks. You do not need to partition virtual disks, because they can be any size you want.
- You must format the partition as NTFS. SVR does not support the FAT file system.
- You cannot use Disk Management to create a volume set, mirror set, or stripe set on the virtual disk. MSCS software and SVR do support RAID software.

Disk space requirements

The space that a virtual disk uses is the same as its capacity. When you create a virtual disk, you specify its capacity in megabytes (MB), and the free space in the pool drops by that number of megabytes. To determine the size of a virtual disk, consider the size of the pool and the number of snapshots you might take. The disk space that 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 create a snapshot, the snapshot does not use any space. The snapshot uses space only when you modify the data stored on the snapshot or its parent disk. The more data you modify, the more space the snapshot uses. The maximum amount of space a snapshot could use as much space as its own capacity (which is the same as the capacity of its parent disk). If you have a virtual disk that is 10 GB, and *all* of the data on the virtual disk changes, a snapshot of the virtual disk could grow to 10 GB.

If you perform incremental backups of your data, you can use the size of your incremental save arrays as an indicator of snapshot size. A snapshot generally uses less space than indicated by the incremental save set because the incremental save set saves the entire file, even if only part of the file has changed.

Backup tools

You can use your backup tools to back up virtual disks and snapshots. If your backup utility performs incremental backups, SVR provides a feature called Incremental Backup Support.

When you back 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. Incremental Backup Support turns off the archive file attributes (or archive bits) of files on a virtual disk after one of its snapshots is backed up.

To enable Incremental Backup Support, you can specify:

- A date and time that Incremental Backup Support 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 prevent the 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 can also create a log file of Incremental Backup Support operations. Refer to the SVR online help for more information.

SVR naming conventions

Every pool, virtual disk, and snapshot must have a unique name. If you name a pool `Pool2`, you cannot create a snapshot named `Pool2`.

No two cluster groups or cluster resources can have the same name. SVR automatically creates cluster groups and resources using the names you choose. Therefore, a pool cannot have the same name as any existing cluster group or cluster resource. For example, if you create a cluster resource called `Res1`, do not create a pool named `Res1`.

To conform to Windows file–naming conventions, avoid the following characters in SVR names:

\ / : * ? " < > |=

Getting Started with Storage Virtual Replicator

3

This chapter describes how to use the Microsoft Management Console (MMC) interface and the command line interface (CLI) to perform SVR tasks:

Using the Microsoft Management Console interface

The MMC interface provides tools that you use to manage local and remote computers from your desktop. SVR's management tool is Replication Manager.

This section describes the following topics:

- Navigating Replication Manager
- Managing remote computers
- Creating a console file

Navigating Replication Manager

To access Replication Manager, select **Start** > **Programs** > **hp OpenView Storage Virtual Replicator** > **Replication Manager**.

The Replication Manager window opens (Figure 1). The structure and format of this window is similar to Windows Explorer.

The scope pane (**A**) provides a tree view of your management tools. You can organize your tools as you would organize files on a disk into folders. For example, you can organize your tools by software vendor and have an HP folder and a Microsoft folder.

When you select an item in the scope pane, the results display in the results pane (**B**), as described in Table 2.

Table 2: Scope and Results pane descriptions

Select	Display
Replication Manager (Local)	All pools (online and offline) on the computer or cluster
Pool	All virtual disks in the pool.
Virtual disk	All snapshots in the virtual disk's family.
Snapshot	Empty

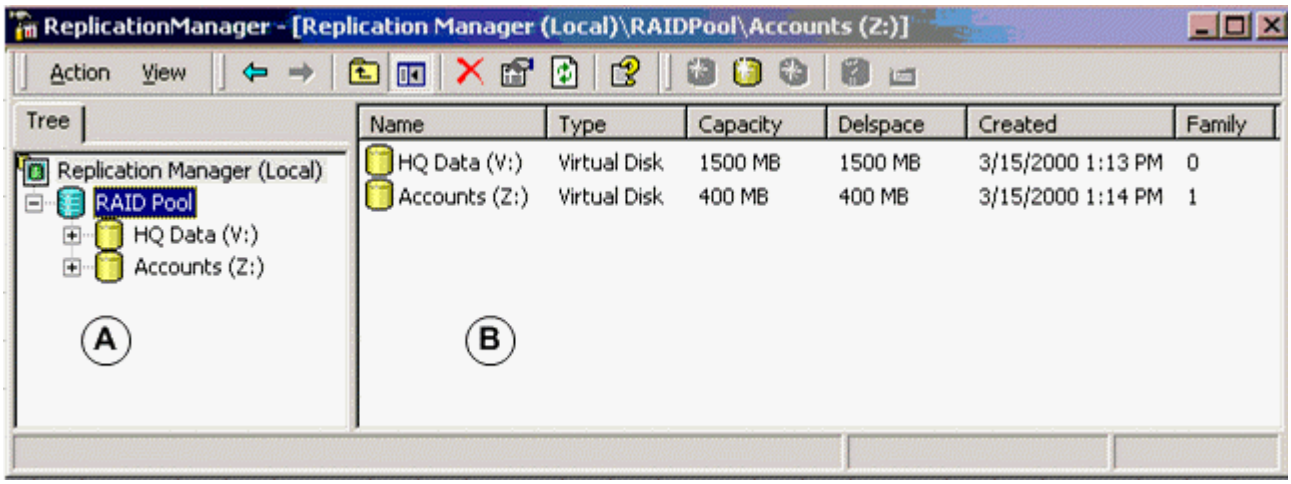

Figure 1: Replication Manager window

You can also use the shortcut menu to perform SVR tasks (Figure 2). Right-click an object in the scope pane to display the shortcut menu (**C**). The shortcut menu actions are also available from the **Action** and **View** menus (**D**) and the toolbar icons (**E**).

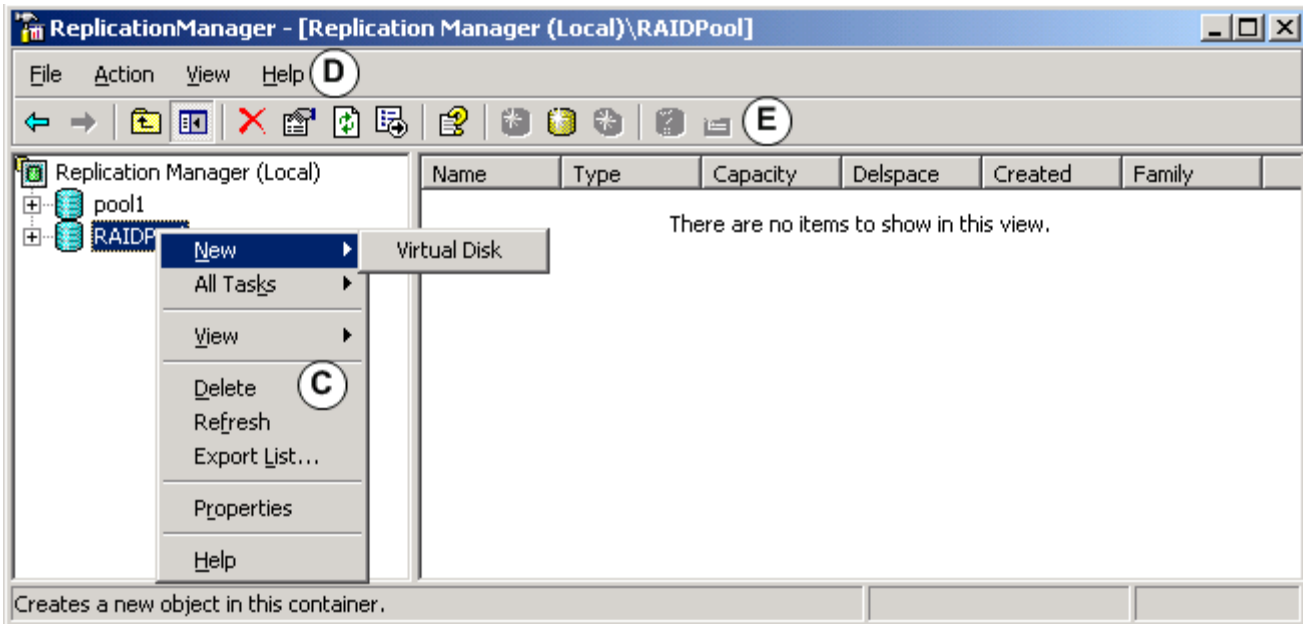

Figure 2: Using the shortcut menu

Managing remote computers

There are three ways you can manage remote computers using SVR:

- Create a temporary connection to a remote stand-alone or cluster computer from the local computer. See "Creating a temporary remote connection."
- Create a console file. See "Creating a console file" on page 34.
- Use the MANAGE command. See "Using the command line interface" on page 37.

Creating a temporary remote connection

If you have SVR installed on a remote computer, you can connect to it from the Replication Manager window that is open on your local computer. The remote connection is temporary; it ends once you close Replication Manager.

To create a temporary remote connection:

1. In the scope pane, right–click **Replication Manager (Local)** and select **Connect to another computer**.

   The Select Computer window opens.

2. Do one of the following:

   - For Windows 2000, enter the name of the computer or cluster in the Name box (Figure 3).
   - For Windows Server 2003, enter the name of the computer or cluster in the Enter object name to select box (Figure 4).

3. Click **OK**.

   You are now connected to the remote computer. The Replication Manager window shows all pools, virtual disks, and snapshots available on the remote computer.

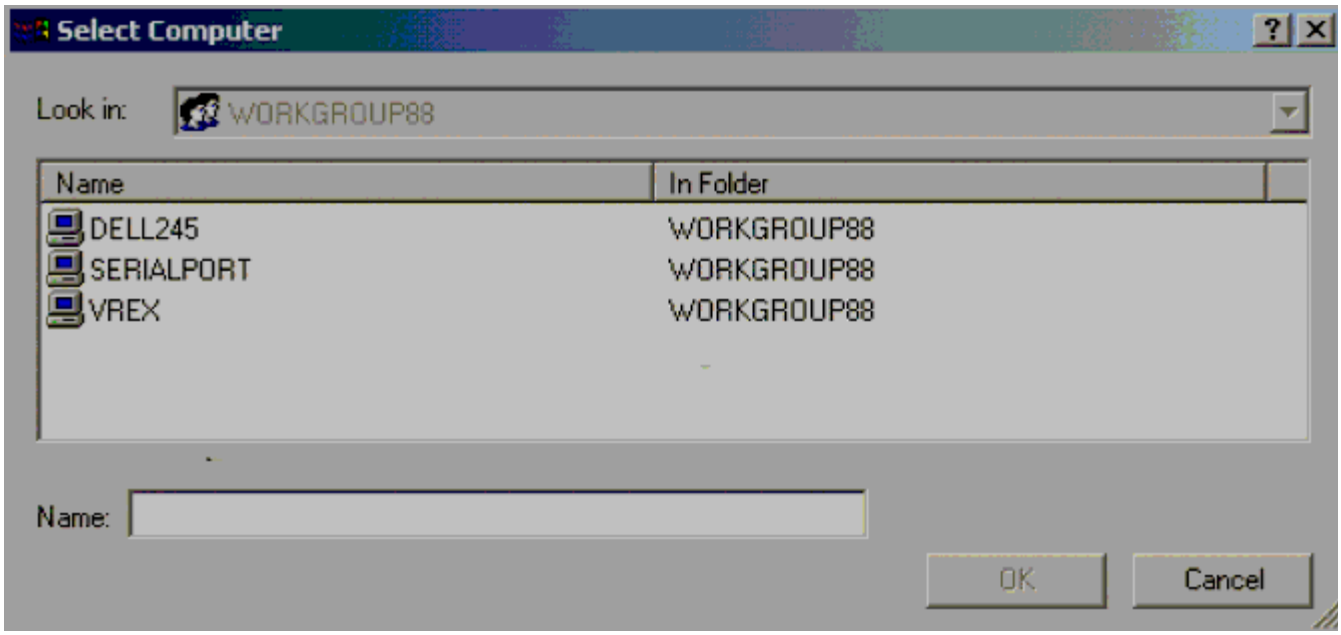

Figure 3: Select Computer window (Windows 2000)

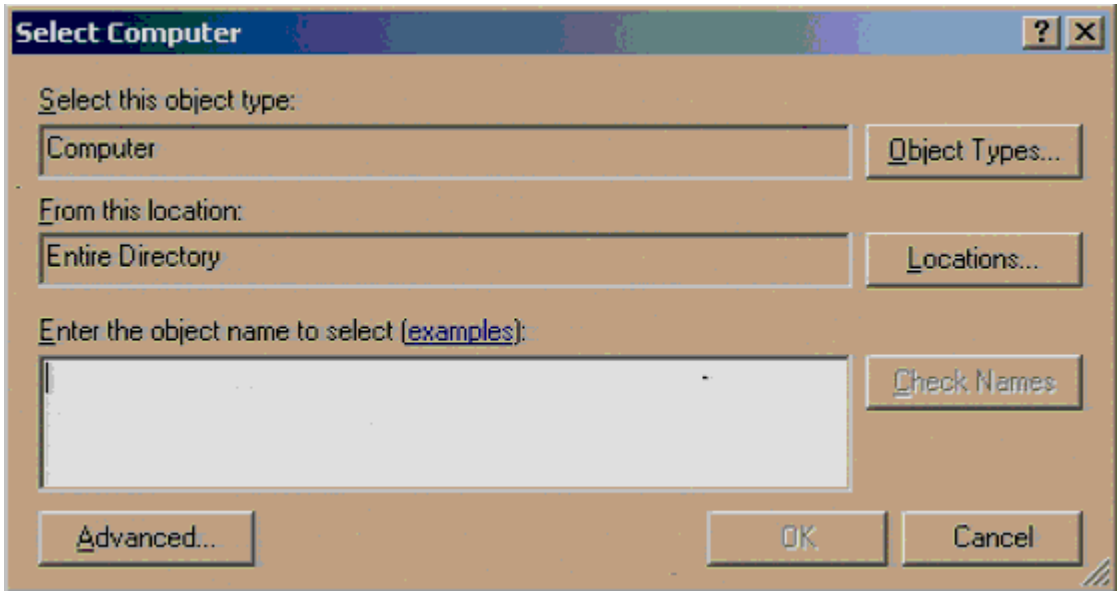

Figure 4: Select Computer window (Windows Server 2003)

Creating a console file

You can create a console file to maintain a permanent connection to a remote computer or to run Replication Manager directly from the MMC interface. A console file is the result of adding a management tool (Replication Manager) to the MMC interface.

Note: In MMC, management tools are called snap-ins.

To create a console file:

1. Select **Start** > **Run**.

   The **Run** dialog box opens.

2. In the **Open** box, enter `mmc` and click **OK**.

   The MMC interface window opens.

3. Select **File > Add/Remove Snap-in**.

   The **Add/Remove Snap-in** window opens (Figure 5).

4. In the **Snap-ins added to** box, select the folder in which you want to add the snap-in.
5. Click **Add**.

   The **Add Standalone Snap-in** window opens.

6. Select **Replication Manager** from the list and click **Add**.
7. Do one of the following (Figure 6):
   - To manage the local stand–alone computer or cluster, select **Local computer**.
   - To manage another stand–alone computer or cluster, select **Another computer** and either enter the name of the computer or cluster or browse for it on the network.
   - Refer to the MMC documentation for information about the **Allow the selected computer** option.
8. Click **Finish**.
9. Click **Close** to close the **Add Standalone Snap-in** window.
10. Click **OK** to close the **Add/Remove Snap-in** window.
11. Select **File > Save As** to save the console file.

    The new console file runs within MMC.

Refer to the MMC documentation for more information about console files.

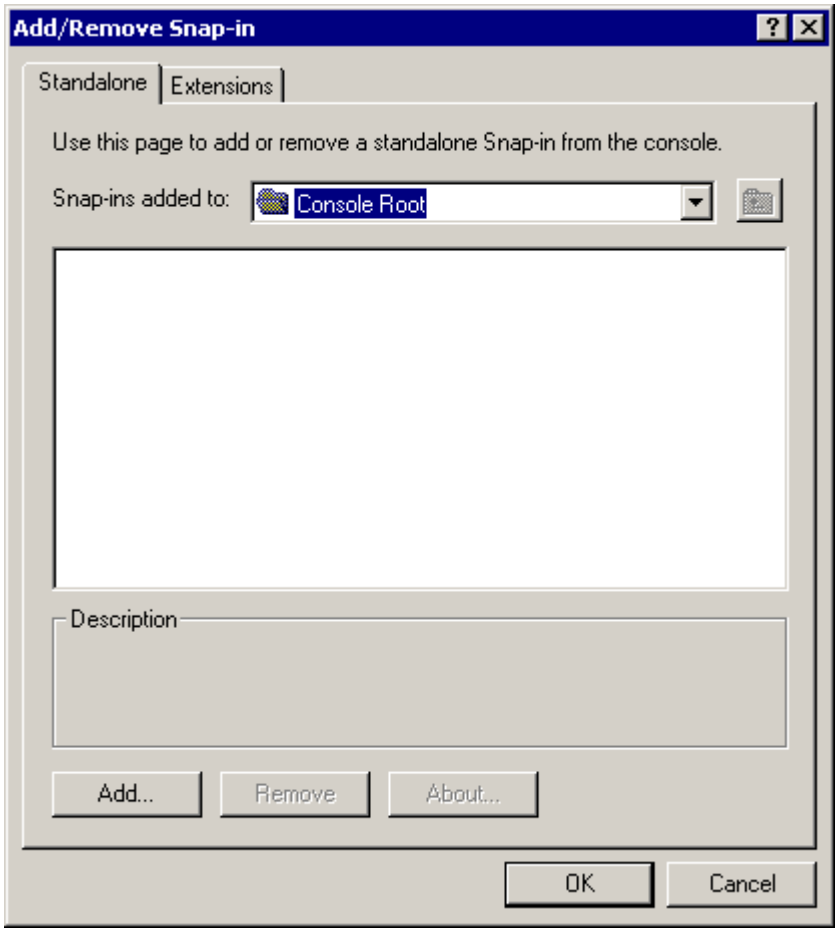

Figure 5: Add/Remove Snap-in window

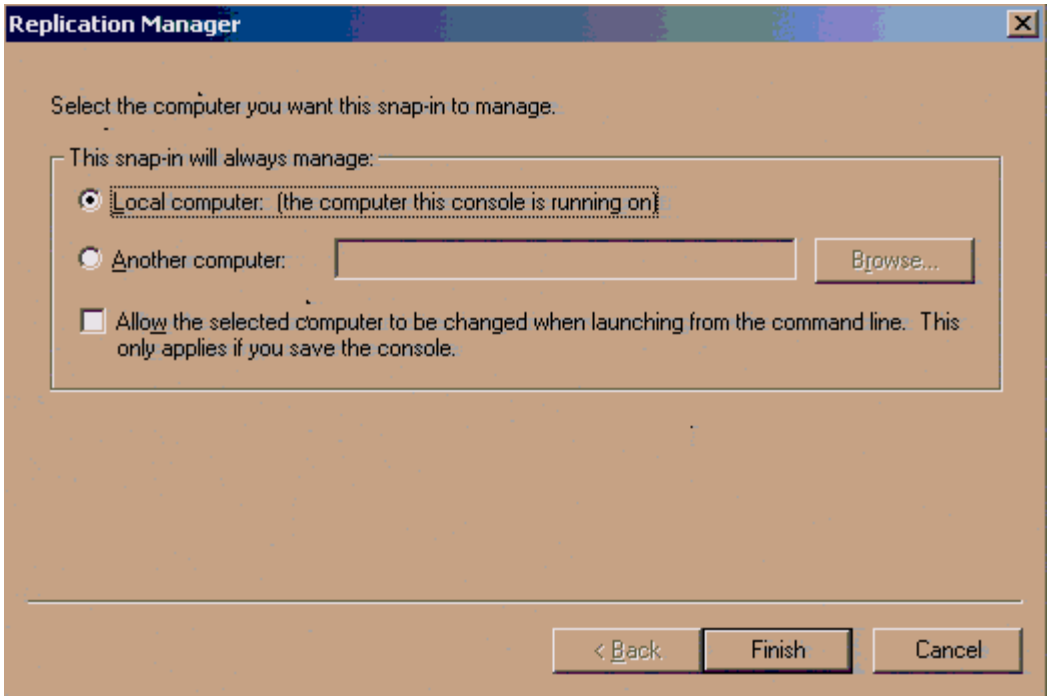

Figure 6: Selecting a computer

Using the command line interface

You can use `SnapMgr`, the SVR command line interface, to:

- Manage pools and the virtual disks and snapshots within those pools.
- Create batch jobs to automate tasks.

You can use the `SnapMgr` commands to manage local and remote computers and clusters. By default, you manage the local stand–alone computer or cluster. Refer to the *HP OpenView Storage Virtual Replicator Command Line Interface Reference Guide* for more information.

Table 3 describes the `SnapMgr` commands.

Table 3: SnapMgr Commands

SnapMgr command	Description
`DRIVES`	Shows which drive letters are available to map to virtual disks and snapshots.
`MANAGE`	Controls the 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

Using SVR with Windows tools

When using SVR with Windows tools, such as Disk Management and Cluster Administrator, consider the following:

- If you have the choice of using SVR or another tool to perform a task, always choose SVR. For example:
  - Use SVR, not Cluster Administrator, to delete a pool in a cluster.
  - Use SVR, not Disk Management, to map drive letters to virtual disks and snapshots.
- Do not use Windows Explorer to monitor free space in a pool. Windows Explorer does not recognize pools.
- Do not use SVR and Disk Management or Logical Disk Manager simultaneously. It may cause unpredictable behavior.
- Close Logical Disk Manager before you use any SVR management tools.
- Do not perform Disk Management tasks, such as mapping drive letters or formatting disks, on disks that belong to SVR pools.
- Do not perform Online Volume Growth or other major disk management operations (such as defragmenting and disk checking) simultaneously.
- Use Device Manager to scan for hardware changes and update disk information before performing any tasks with Disk Management.

  To open Device Manager, right–click **My Computer** and select **Properties > Hardware**.

Managing security and privileges

By default in Windows, the Domain Admins global group is a member of the local Administrators group on every computer. Any user in the Domain Admins group can use SVR to manage any computer or cluster in the domain.

Creating a group

To enable users to manage SVR without having the privileges of the Domain Admins group:

1. Create the SWVR Admins group on the domain controller. You must have Domain Admins privileges to create this group.
2. Add users to the SWVR Admins group.
3. Add the SWVR Admins group to the local Administrators group on each computer you want users to manage, including every node in a cluster.

To use Online Volume Growth, you must have Full Control permission at the top level of the volume you want to grow. By default, administrators have Full Control permission for NTFS volumes on Windows systems. You do not need to modify this permission unless it has been explicitly denied.

Managing other domains

You can manage a computer in another domain if:

- The domains are compatible (trust relationship).
- Your user account belongs to the local Administrators group on every computer you want to manage, including every node in a cluster.

Using cluster resources

Follow these guidelines when using SVR on a Microsoft Cluster Server (MSCS) system:

- When you create pools in a cluster, SVR automatically creates all required groups and resources. It also creates the correct dependencies, possible owners, and default properties so that pools can fail over within the cluster.
- For each pool, SVR creates a cluster group named *poolname* Group (where *poolname* is the name you assigned to the pool). It also creates a resource for the pool itself. This resource has the same name as the pool and has the SCE Pool type.
- Do not use Cluster Administrator to create or rename any of these resources. If you do, 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 the node that owns the pool resource.
- Maintain a backup copy of the Cluster Registry in case you need to perform a system recovery.
- By default, all nodes in the cluster are potential owners of the pool resource. If you want to restrict the owners of the pool resource, use Cluster Administrator. Refer to the Microsoft Cluster Administrator documentation for instructions.
- After failover, SVR cluster groups do not fail back automatically. To configure a group for failback, use Cluster Administrator. See “Creating pools in a cluster” on page 46 for more information.

Managing Pools

4

This chapter describes how to create and manage pools:

Creating pools

When creating pools, consider the following:

- You create pools from unformatted (raw) disks.
- The number of pools is limited only by the number of storage units that are available.
- After you create a pool, SVR marks the storage units as offline so users cannot access them.

To create a pool:

1. From the Replication Manager window, right–click **Replication Manager (Local)** and select **New > Pool** (Figure 7).

   The New Pool Wizard window opens.

2. Click **Next** to continue.

   The Storage Units window opens (Figure 8).

3. Select the disks you want to use for the pool in the Available storage units box and then click the right arrow.

   In this example, you select Disk 2 and Disk 4. Disk 2 and Disk 4 move to the Selected storage units box (Figure 9).

   The Total capacity field shows the total capacity of the selected disks. The pool capacity can be a maximum of 10 percent less than the total capacity, because the internal pool configuration data uses disk space.

4. Click **Next**.

   The Pool Information window opens.

5. Enter a name for the pool in the Pool name box.

   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. In this example, you enter `RAID Pool` (Figure 10)

6. Select a segment size for the pool or accept the default size of 128 KB under Segment size.

   The segment size defines the largest virtual disk you can create from the pool and the largest storage unit you can add to the pool. It also defines the smallest unit you can copy out to a snapshot when changes occur on a parent disk.

Note: Once you select a segment size and complete the pool creation, you cannot change the segment size.

---

7. Click **Next**.

   The Completing the New Pool Wizard window opens, indicating if the pool creation was successful. If pool creation fails, this window displays an explanation for the failure.

8. Click **Finish**.

   The Replication Manager window displays the new pool (Figure 11).

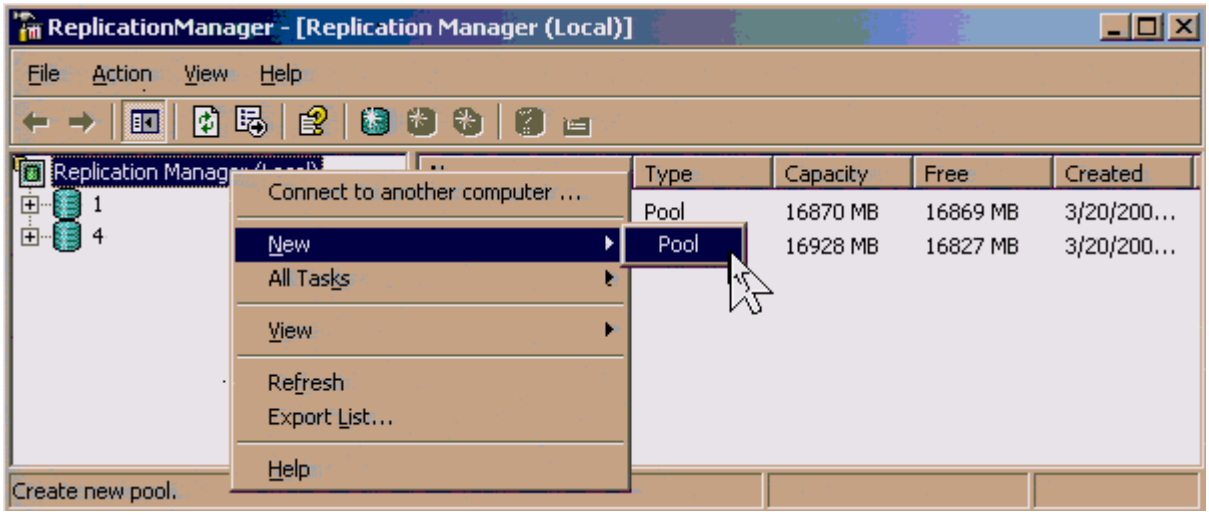

Figure 7: Replication Manager window

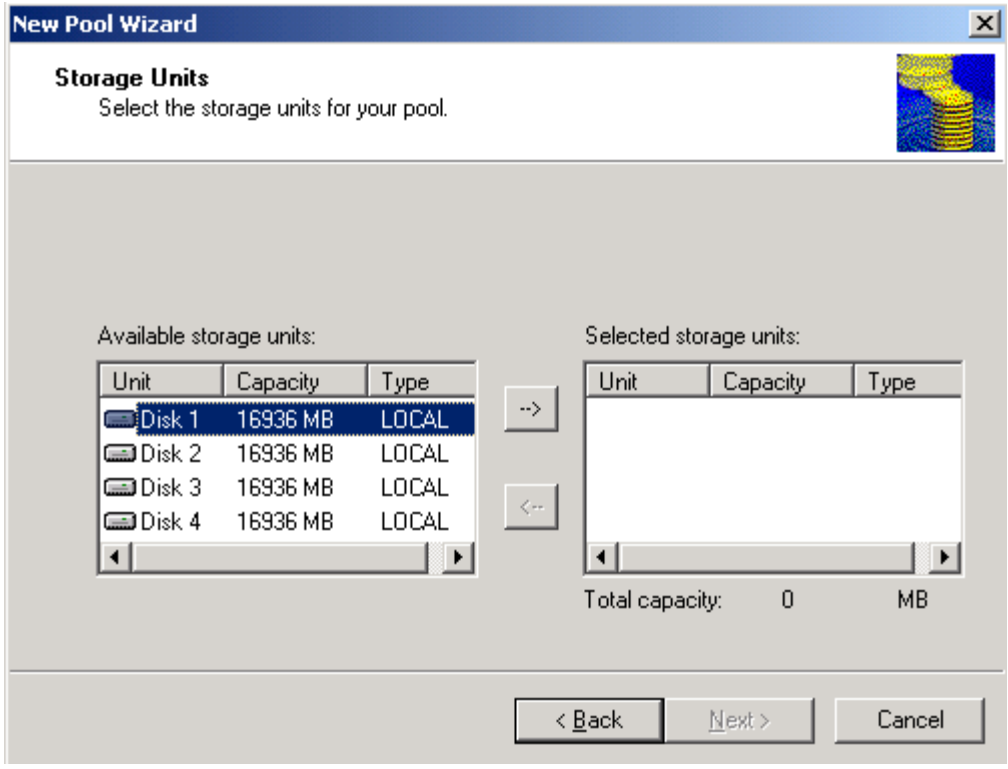

Figure 8: Storage Units window

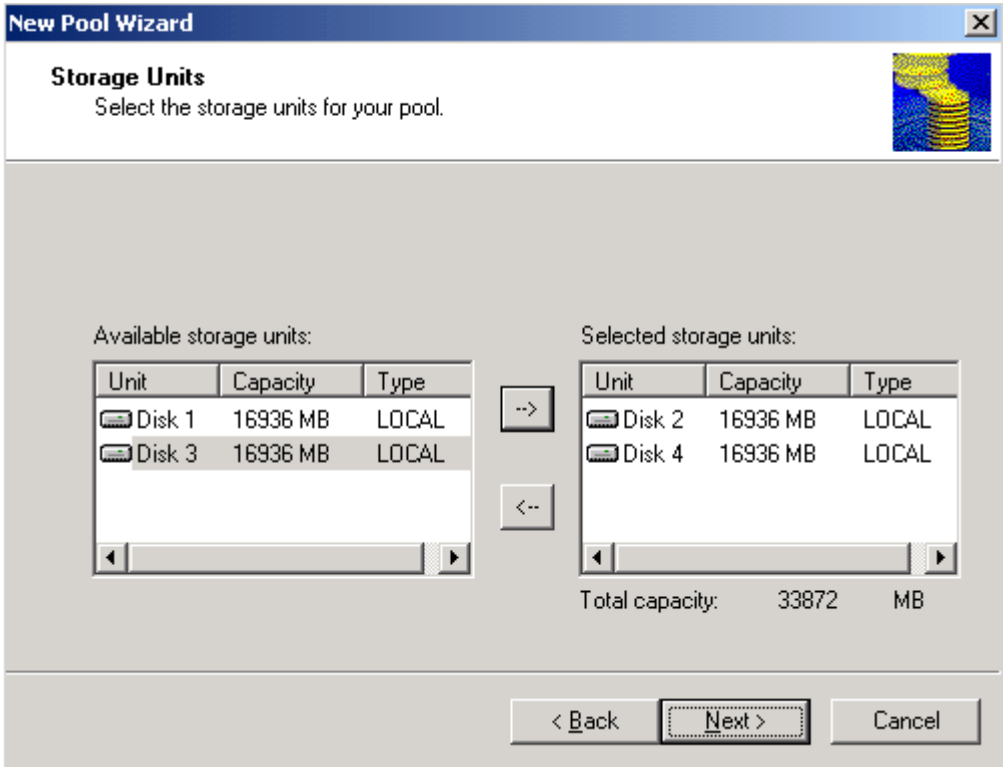

Figure 9: Storage Units window with selected disks

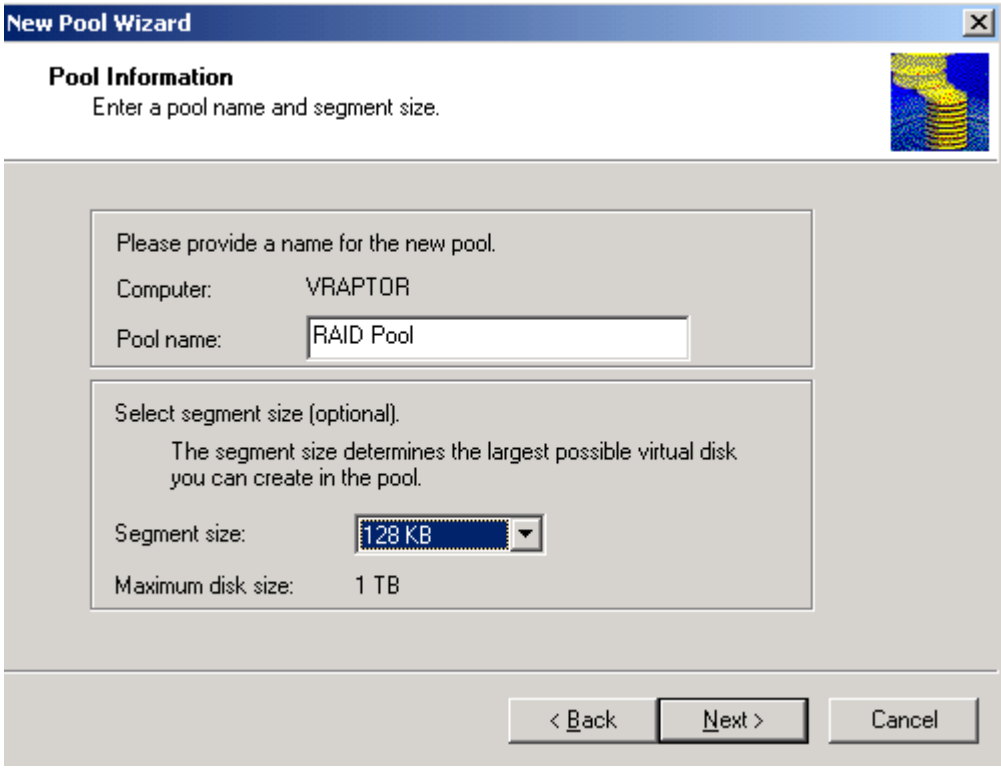

Figure 10: Pool Information window

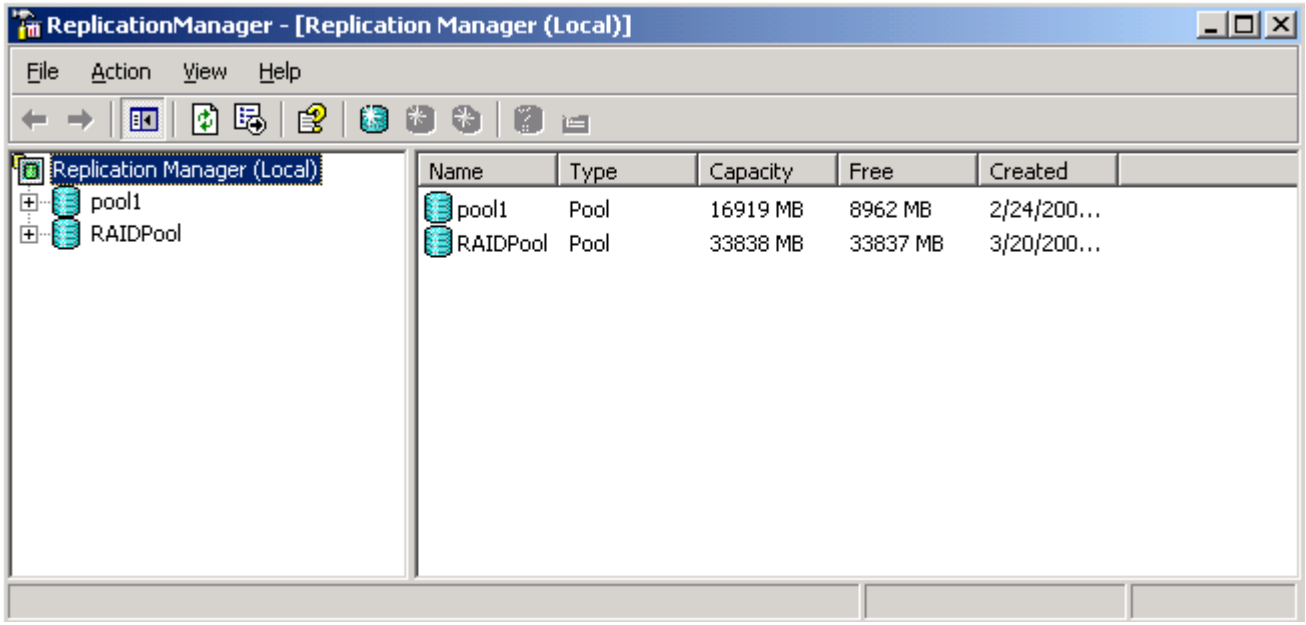

Figure 11: Replication Manager window with new pool

Creating pools in a cluster

When you create pools in a cluster, the pool wizard:

- Creates a cluster group named *poolname* Group (where *poolname* is the name for the pool).
- Creates a resource type called SCE Pool for the pool. The resource has the same name as the pool.
- Brings the pool group and pool resource online.

Caution: Do not use 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*, you can rename the pool *group*. You can have more than one pool in the group and you can move pools to 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 causes the pool to start up before the application or file share.

Specifying owners for pool resources

When you create pools in a cluster, all nodes in the cluster are possible owners of the pool resource. However, you can specify which nodes own the resource.

Open the Properties window for the pool resource using Cluster Administrator. You can add or remove nodes from the list of nodes under Possible owners. Refer to the Microsoft Cluster Administrator documentation for instructions.

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

Using storage units in pools

This section describes how to use storage units in pools:

- Adding storage units to a pool
- Importing storage units into a pool
- Removing storage units from a pool

Adding storage units to a pool

After you create a pool, you can add storage units to it at any time, even if users are accessing the pool's virtual disks and snapshots.

When adding storage units to a pool, consider the following:

- A pool may contain up to 12 storage units.

Note: HP recommends that you use a maximum of 10 storage units in a pool in case you need to migrate data to the pool and import additional storage units.

- Storage units can be standard single-spindle disks or controller-based, fault-tolerant disk sets. They can have different capacities and can be from different manufacturers.
- The maximum physical disk capacity is 2 TB.
- All storage units should have the same read-write, redundancy, and failure characteristics. For example, all units should be standard disks, or RAID 5 storage sets, or hardware mirror sets.
- In a cluster, use the disks on the shared storage bus only. Do not use local disks.
- After you add storage units to a pool, they are automatically marked as offline so that users cannot access them.

Importing storage units into a pool

You can import storage units with active data into a pool without disruption or downtime. When you import a partitioned storage unit, SVR automatically creates a virtual volume for each partition in the unit. The applications that used data on the physical unit can continue to access the data on the virtual disk.

If the physical unit has space available, you can add it to the pool's free space after importing the unit. You can then use this additional space to store other data.

When importing storage units into a pool, consider the following:

- Ensure that the pool's free space is equal to or greater than 1% of the storage unit's capacity.
- SVR does not support the FAT file system. If the storage unit you want to import has FAT partitions, you must convert them to NTFS before you import the unit.
- In a cluster, the storage unit you want to import must be a shared disk and 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 provided by the drive. To use this space, you must add the drive to the pool as a separate LUN.

  You can add a single drive to the storage system if it does not interfere with the storage units that belong to the pool.

Refer to the SVR online help for more information.

Removing storage units from a pool

You can 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 virtual disks are using that unit. Since you cannot identify which storage units are being used, you should only remove a storage unit that you have just added to a pool.

Refer to the SVR online help for more information.

Managing pool free space

This section describes how to manage free space for pools:

- Defining pool free space
- Monitoring pool free space
- Using Windows Explorer
- Calculating Delspace

Defining pool free space

Pool free space is disk space that is not allocated to virtual disks, snapshots, or internal configuration data. Free space prevents the pool from becoming full, which may cause failed writes and data loss.

If the pool capacity is less than the combined capacity of its virtual disks and snapshots, the pool can become full. If a pool runs out of disk space, you must do one of the following:

- Add a storage unit to the pool.
- Delete one or more snapshots or virtual disks.

Caution: Do not delete files to create free space. Deleting files causes segments of data to be copied out to snapshots, which uses more space in the pool and may cause data loss.

Monitoring pool free space

To manage pools effectively, monitor their free space regularly. It is important to know how much space is gained when you delete virtual disks and snapshots.

Note: Always use SVR to monitor pool free space.

To monitor free space, you can:

- Check the Free column for a pool in the Replication Manager window (Figure 12).
- Check the system event log for SVR warnings and errors about pool free space:
  - When the free space falls below its threshold (the default is 30 percent), SVR writes a warning to the log every five minutes.
  - When the free space falls below five percent, SVR writes an error to the log every five minutes.

  See "SVR driver events" on page 149 for more information.
- Use the SVR Lifeguard service to control monitoring activities. See "Setting SVR Lifeguard policies" on page 134 for more information.

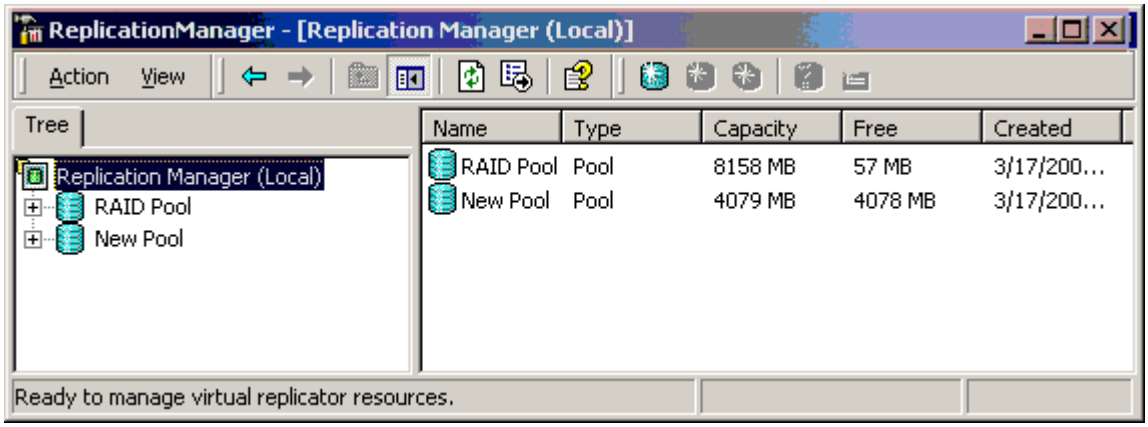

Figure 12: Monitoring pool free space

Using Windows Explorer

Do not use Windows Explorer to monitor free space because it does not recognize pools. For example, if the RAID Pool pool shown in Figure 12 contained one virtual disk and two snapshots, the following information displays in the Windows Explorer Properties window:

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

In Windows Explorer, the combined free space on the three disks is 2,650 MB. However, the actual free space in the pool, as shown in Figure 12, is only 57 MB.

Deleting files has different effects in Windows Explorer and Replication Manager. For example, if you delete files on the `Z:\` drive (VirtualDisk), the amount of free space increases in Windows Explorer. If you delete files in Replication Manager, the free space does not change because deleting files does not gain pool space.

Calculating Delspace

The Delspace of a virtual disk or snapshot is the amount of pool free space gained when you delete a virtual disk or snapshot. This section describes the following topics:

- Virtual disk Delspace
- Snapshot Delspace

Virtual disk Delspace

To view the Delspace of all virtual disks in a pool, select the pool in the scope pane of the Replication Manager window. The results pane shows the details for all virtual disks in the pool (Figure 13).

The Delspace of each virtual disk is equal to zero or the capacity of the virtual disk. If a virtual disk contains snapshots, the Delspace is zero because you cannot delete the virtual disk if it contains snapshots. After you delete the snapshots, you can delete the virtual disk. The Delspace of the virtual disk then changes from zero to the capacity of the virtual disk.

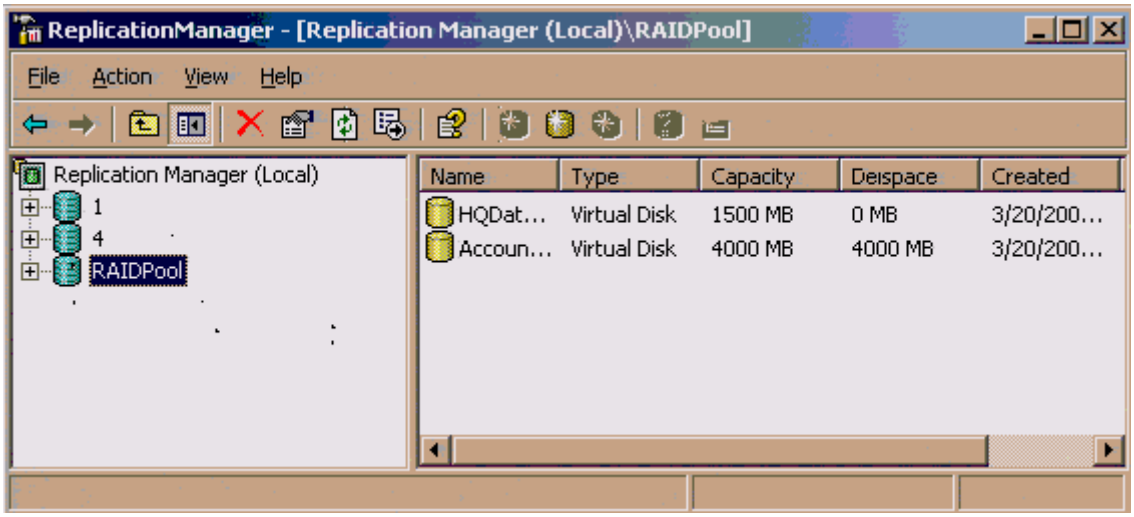

Figure 13: Viewing the Delspace of virtual disks

Snapshot Delspace

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

To view the Delspace of a snapshot, select the virtual disk for the snapshot's family in the scope pane of the Replication Manager window. The results pane shows the details for all snapshots in the family.

In Figure 14, the Delspace for the *Accounts Mon* snapshot is 450 MB. If you delete *Accounts Mon*, the free space in the pool increases by 450 MB.

When deleting snapshots to increase free space, consider the following:

- If there is only one snapshot in a family, the Delspace is the amount of space that the snapshot is using in the pool.
- If there are two or more snapshots in a family, the snapshots can share disk space. You cannot determine how much space each snapshot uses.
- If you delete two snapshots in a family, you can gain more space than the sum of their Delspaces.

  For example, in Figure 14, if you delete both snapshots, you gain 1,000 MB of space even though the sum of the Delspaces is 624 MB. The difference between the free space you gain (1,000 MB) and the sum of the Delspaces (624 MB) is the amount of disk space that *Accounts Mon* shared with *Accounts Tue* (376 MB). If you delete *Accounts Mon*, the Delspace for *Accounts Tue* increases by 376 MB (174 MB to 550 MB), which is the amount of disk space that was shared between the snapshots and is now used exclusively by *Accounts Tue*.

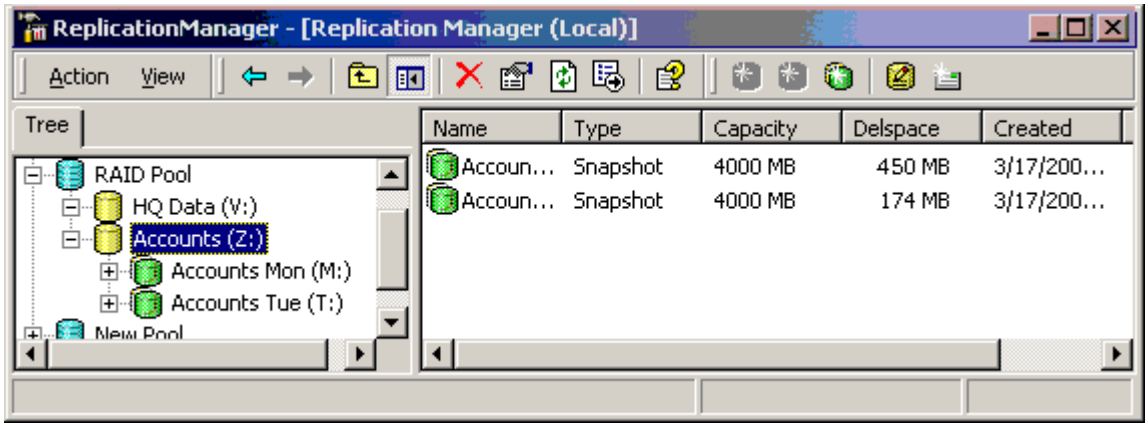

Figure 14: Viewing the Delspace of snapshots

Viewing pool information

You can view summary information for pools using Replication Manager and detailed information using the **Properties** window. This section describes the following topics:

- Using Replication Manager
- Using the Properties window

Using Replication Manager

The results pane in the Replication Manager window displays the following information for all pools on a stand–alone computer or cluster:

- **Name**—The name of the pool.
- **Type**—The value is `Pool`.
- **Capacity**—The capacity of the pool, rounded up to the nearest megabyte.
- **Free**—The amount of available pool free space, rounded down to the nearest megabyte.
- **Created**—The date on which the pool was created.
- **Owner**—The name of the cluster node that owns the pool resource. This information is not displayed for a stand–alone computer.

Note: In clusters, Replication Manager does not show pools that are offline.

Using the Properties window

To view specific information for a pool, right–click the pool and select **Properties**. The Properties window opens (Figure 15).

The Properties window contains four tabs:

- **General**—Displays general pool information. See page 54 for details.
- **Storage Units**—Displays information about the storage units in the pool. See page 55 for details.
- **Statistics**—Displays statistical information about the pool. See page 56 for details.
- **Policy Editor**—See “Using Policy Editor” on page 137 for more information.

The General tab displays the following information:

- **Type**—The value is `Pool`.
- **Status**—The status of the pool. Values are `Online`, `Online Pending`, and `Offline`.
- **Capacity**—The capacity of the pool.
- **Free space**—The amount of free space available in the pool.
- **Segment size**—The segment size of the pool. A segment is the smallest unit of data that can be moved during a copy-out operation.
- **Created**—The date on which the pool was created.
- **Modified**—The date on which the pool last modified (that is, the date on which a storage unit was added to the pool).

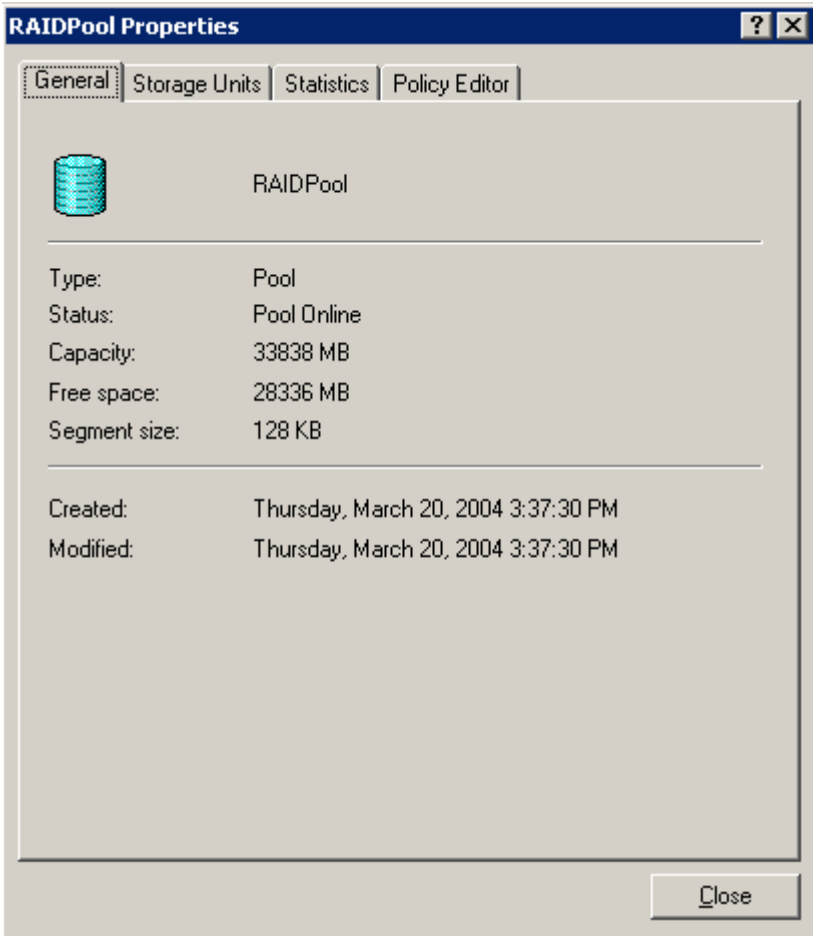

Figure 15: Properties window

The Storage Units tab (Figure 16) displays the following information:

- **Unit**—The disk number of the unit. In a cluster, the disk number is from the node that owns the pool resource. The disk number may change when you restart the computer or when the pool resource fails over in the cluster.
- **Capacity**—The capacity of the storage unit.
- **Type**—The type of storage unit. The value is `Shared` for a disk in a cluster that is on a shared storage bus and can be seen by every node in the cluster. The value is `Local` for a physical disk that is attached to a stand–alone computer.

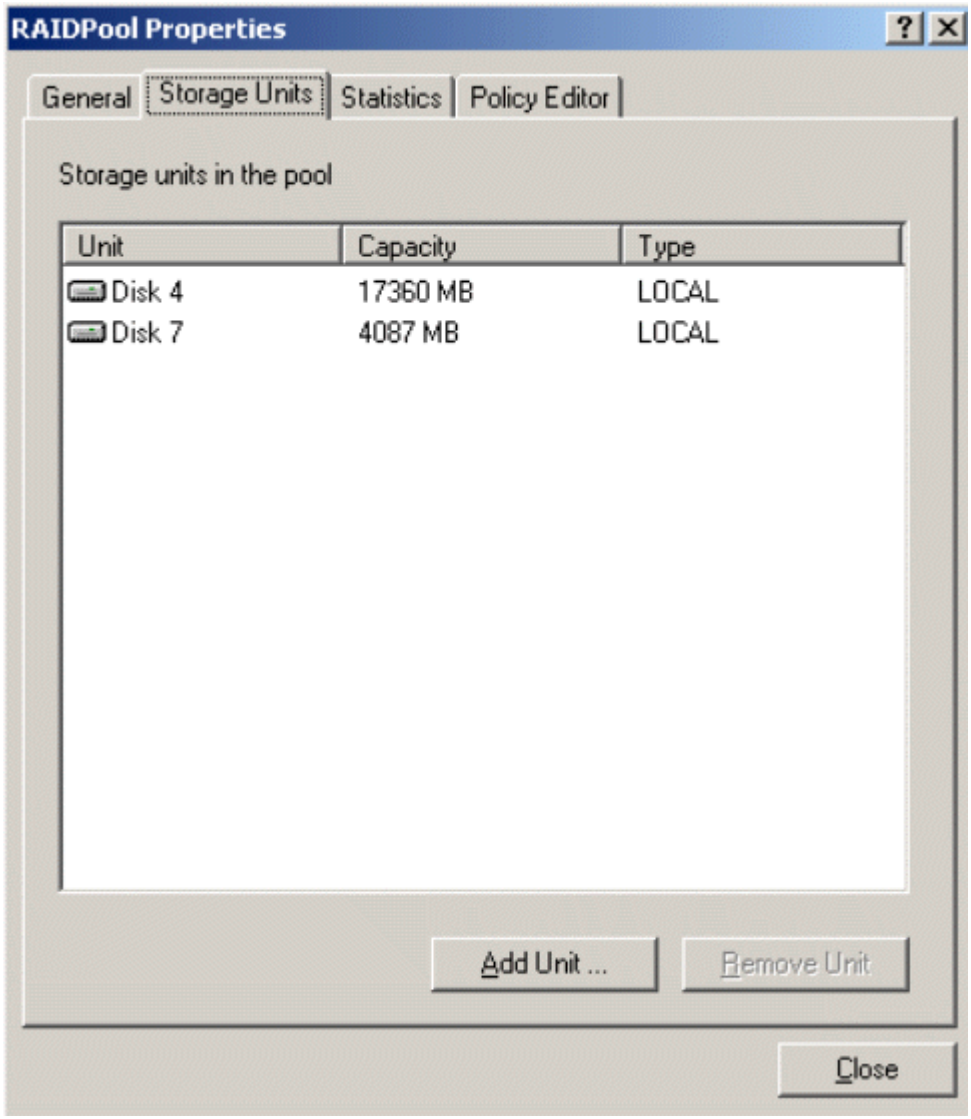

Figure 16: Storage Units tab

The Statistics tab (Figure 17) displays the following information:

- **Reads received**—The number of read requests the pool received.
- **Reads issued to disk**—The number of read requests the pool issued to disk, including read requests caused by copy-outs and split reads.
- **Split reads**—The total number of read requests that were divided into two or more requests because the original request crossed a segment boundary, and the next segment was not contiguous with the current segment.
- **Writes received**—The number of write requests the pool received.
- **Writes issued to disk**—The number of write requests the pool issued to disk, including write requests caused by copy-outs and split writes.
- **Split writes**—The total number of write requests that were divided into two or more requests because the original request 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.

These counters are reset to zero when the computer restarts or, in a cluster, when the pool fails over.

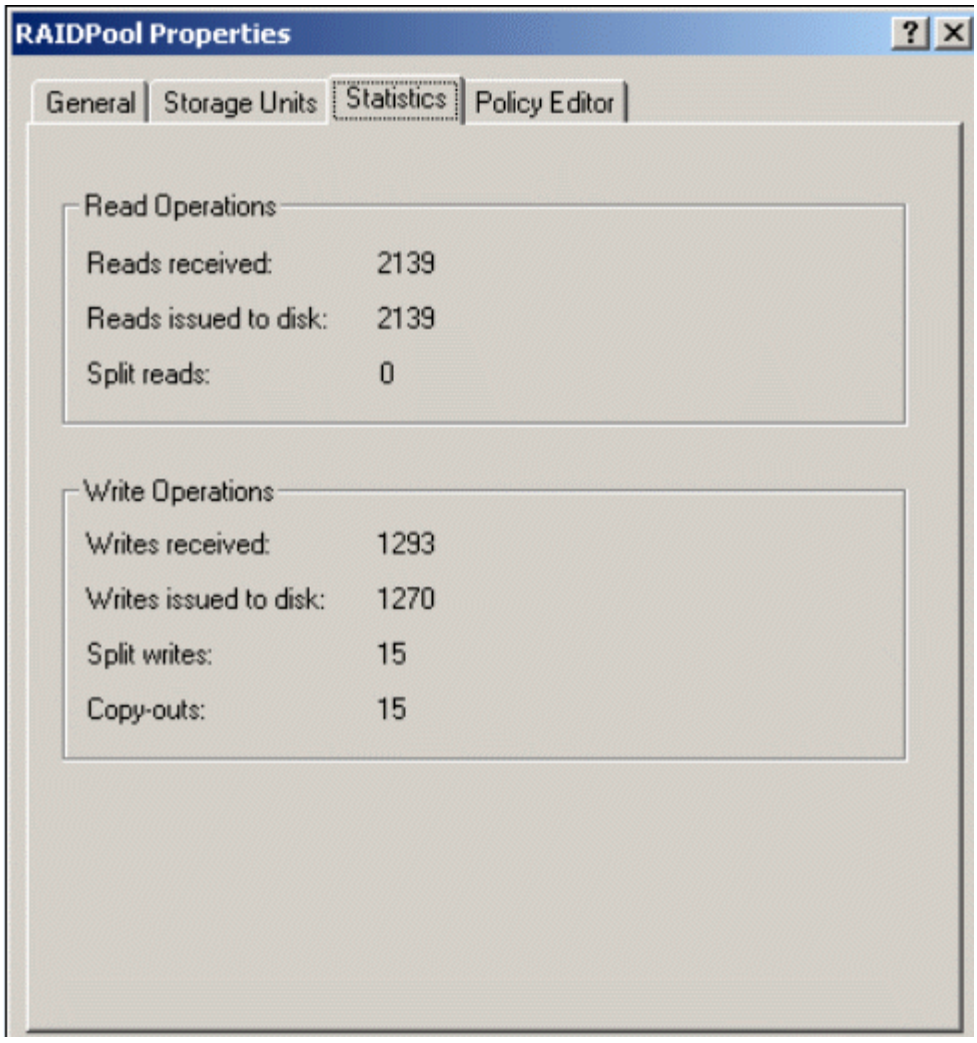

Figure 17: Statistics tab

Deleting pools

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

Deleting a pool automatically deletes the pool resource and the pool group, if the group contains no other resources.

Refer to the SVR online help for more information.

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

Managing Virtual Disks

5

This chapter describes how to create and manage virtual disks:

Note: If there are two ways to perform a task–using Replication Manager or another tool–always use Replication Manager. For example, use Replication Manager to map drive letters to virtual disks, not Disk Management.

Creating virtual disks

When creating virtual disks, consider the following:

- You can only create virtual disks after you create a pool.
- You can create a maximum of 12 virtual disks for each pool.
- Virtual disks use space from the entire pool. You cannot specify which storage unit a virtual disk uses.
- Virtual disks function the same as physical disks, enabling you to meet the disk space requirements of your users and applications.
- The virtual disk capacity must be a minimum of 10 MB and a maximum of 2 TB. The capacity is determined by the pool's free space, segment size, and policies.
- There are no cluster resources for virtual disks or snapshots. The virtual disks and snapshots in a pool are displayed on the node that owns the pool resource.

Note: If the drive letter you map to the virtual disk is not available when the pool fails over (because the drive letter was assigned to a partition on the second node), SVR dismounts the drive letter and reassigns it to the failed over volume.

To create a virtual disk:

1. In the scope pane, right–click the pool and select **New** > **Virtual Disk**.
   The New Virtual Disk Wizard window opens.
2. Click **Next**.
   The Virtual Disk Information window opens (Figure 18).

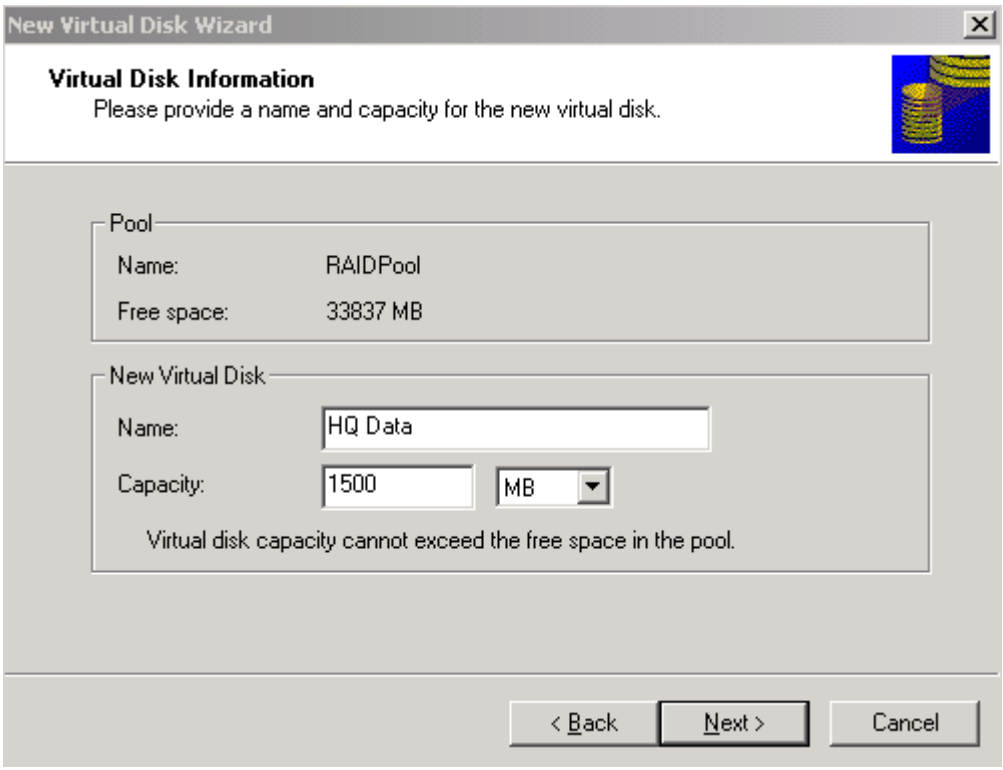

Figure 18: Virtual Disk Information window

3. Enter the name of the virtual disk in the Name box and click **Next**.
4. Enter the size of the virtual disk in the Capacity box and click **Next**.
   The Drive Letter Assignment window opens (Figure 19).

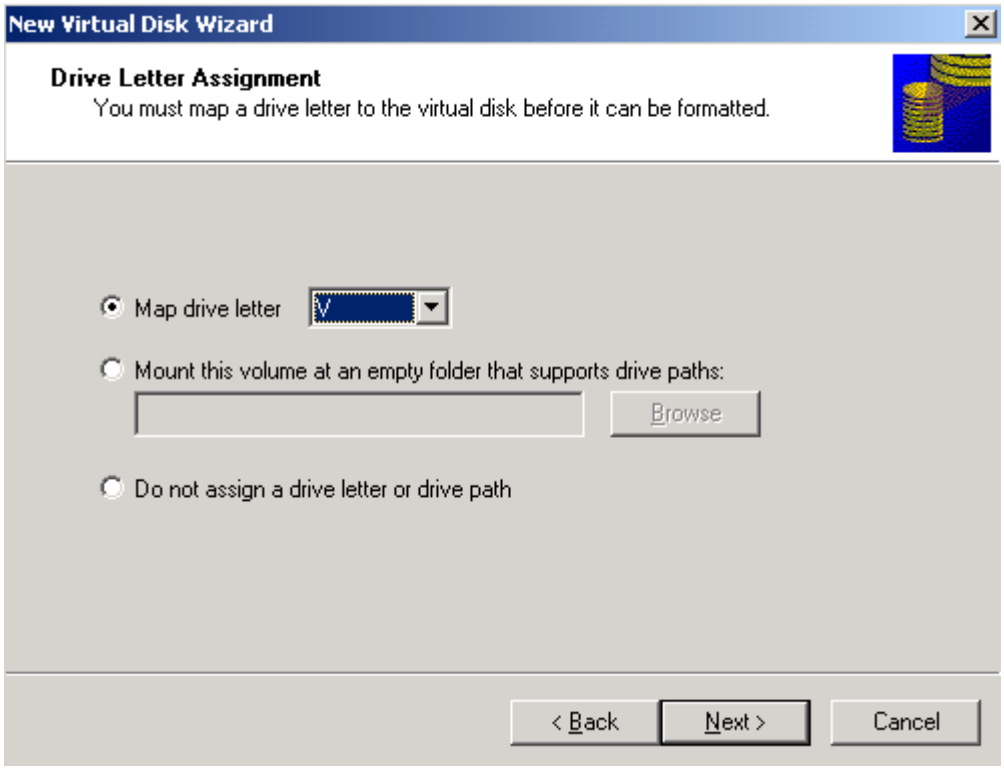

Figure 19: Mapping a drive letter

5. Do one of the following:
   - Select **Map drive letter** and choose the drive letter you want to map to the virtual disk (Figure 19).

     Choose a drive letter near the end of the alphabet to ensure it is not allocated to another disk. See "Mapping drive letters" on page 68 for more information.
   - Select **Mount this volume at an empty folder that supports drive paths**. Click **Browse** and choose an empty folder on an NTFS volume on which to mount the drive (Figure 20). Click **OK**.

     Refer to the SVR online help for best practices in creating mount points in a clustered environment.

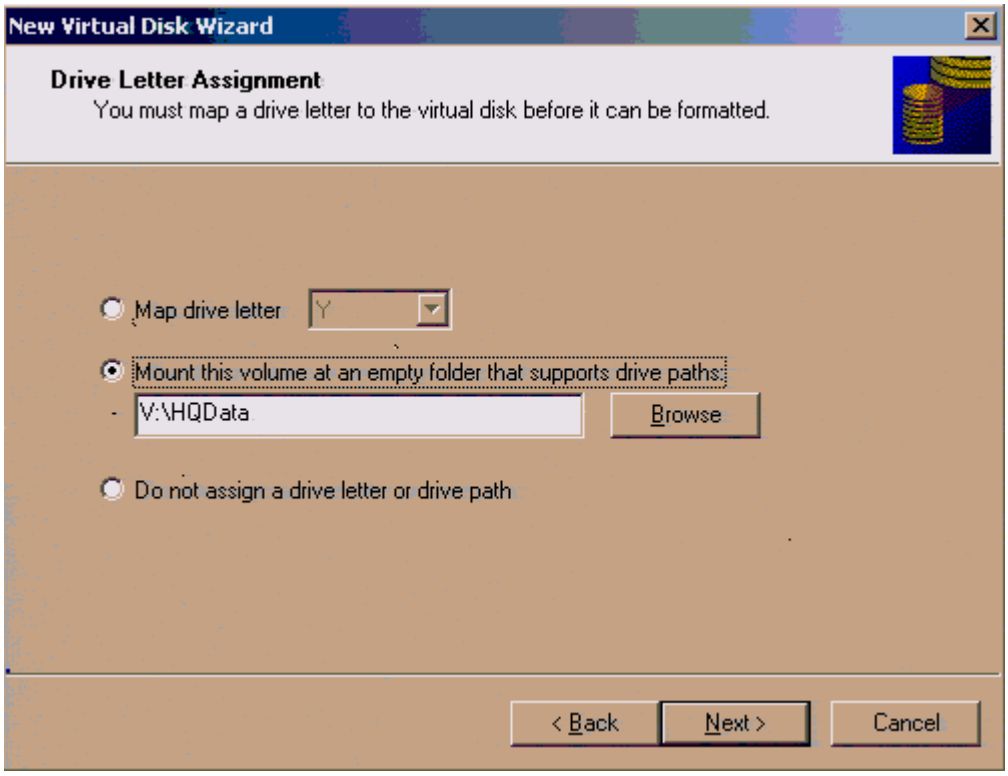

Figure 20: Mounting a volume

6. Click **Next**.

   The window that opens depends on the option you selected in step 5.

   - If you selected the mount option, the Completing the New Virtual Disk Wizard window opens, indicating if the virtual disk creation was successful. If virtual disk creation fails, this window displays an explanation for the failure.
   - If you selected the map option, the Format Virtual Disk window opens (Figure 21). Select **Yes** to format the virtual disk. See "Formatting virtual disks" on page 69 for more information.

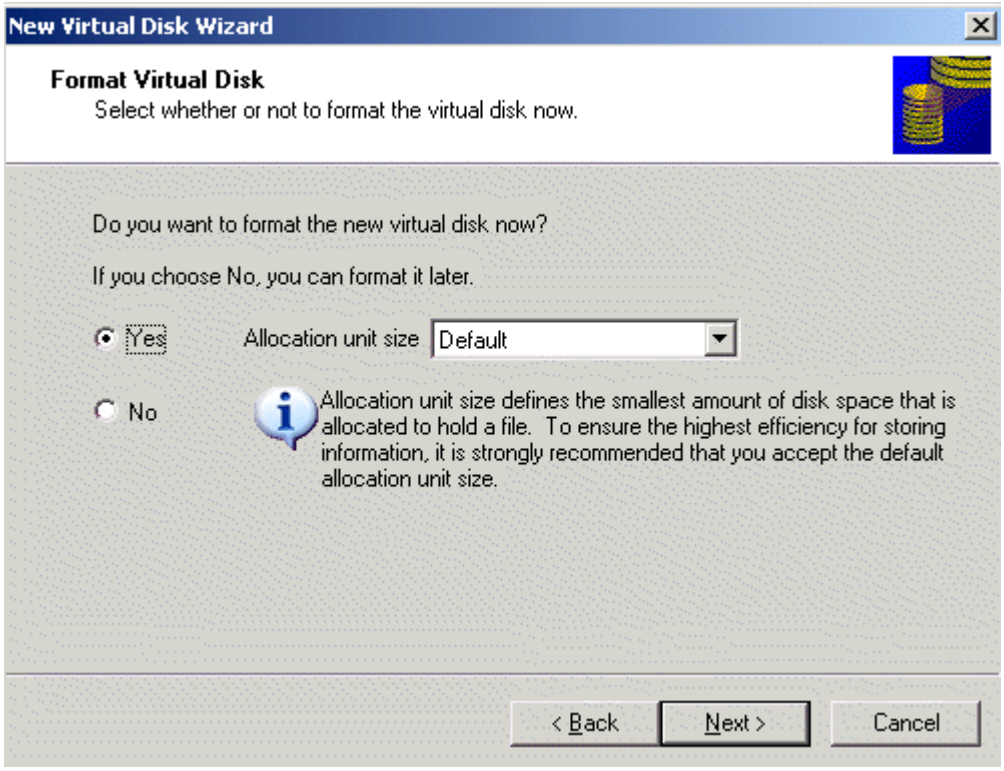

Figure 21: Formatting a virtual disk

7. After formatting the disk, click **Finish** on the Completing the New Virtual Disk Wizard window.
8. In the scope pane, select the Replication Manager (Local) folder (Figure 22).

   The new virtual disk (HQData) is displayed under RAIDPool in the scope pane. In the results pane, the pool's free space is now 32,337 MB. The free space decreased by 1,500 MB, which is the capacity of the virtual disk you just created.

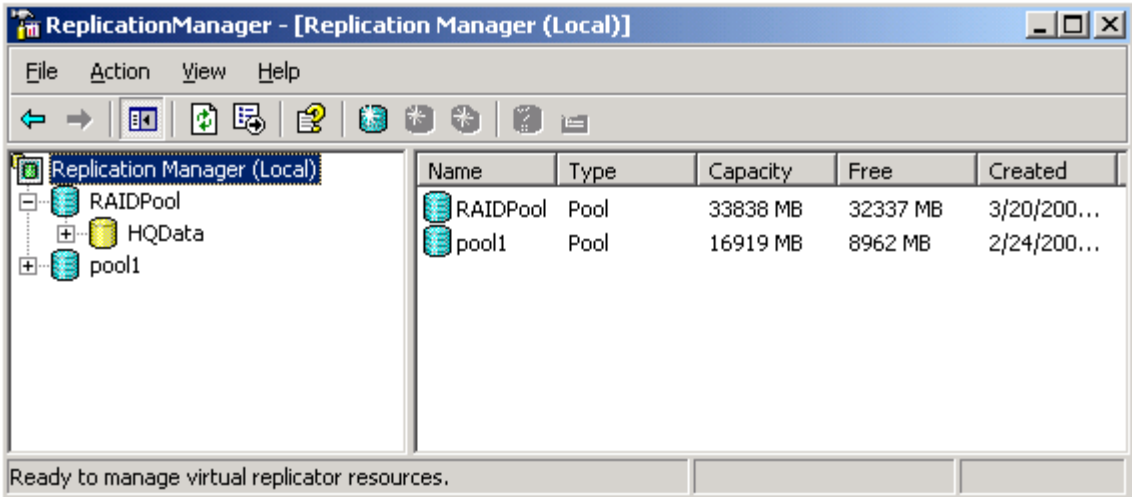

Figure 22: Viewing the new virtual disk

9. From the Windows desktop, open My Computer.
   The new virtual disk is displayed with drive letter (v) and volume label (HQ Data).
10. Copy the Personal folder and the `myfile.txt` file to the new virtual disk (Figure 23).

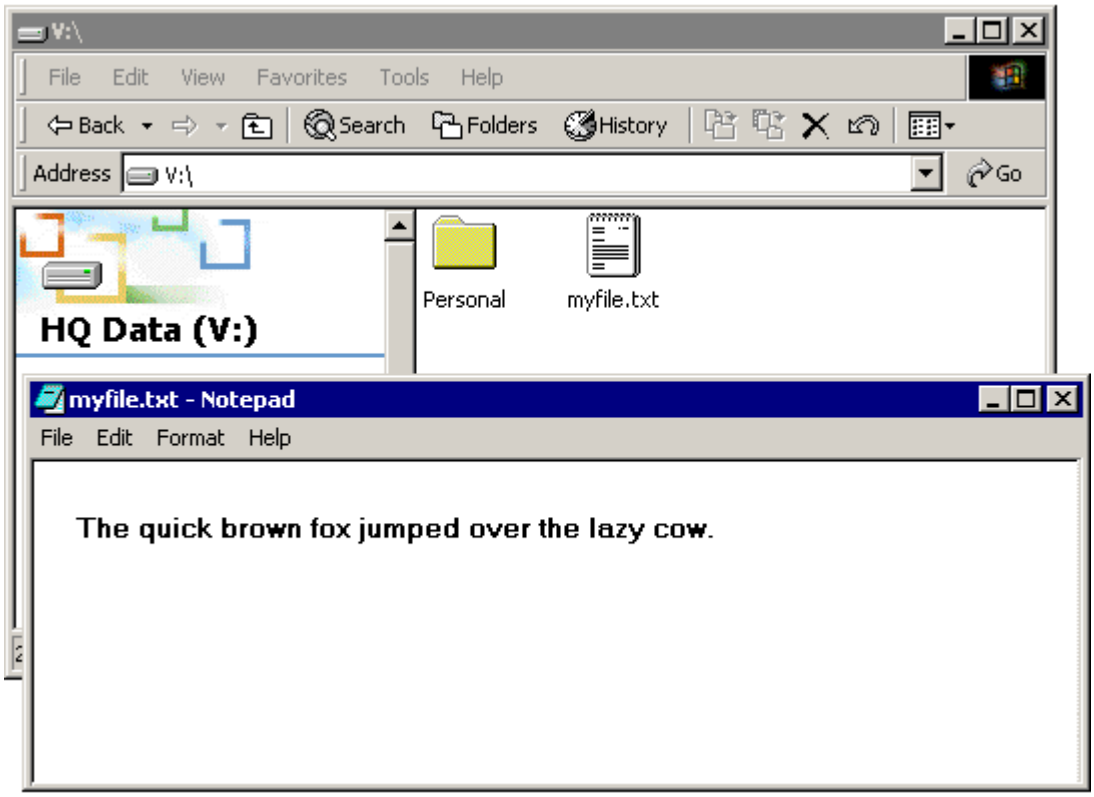

Figure 23: Copying a folder and file to the virtual disk

11. Click the virtual disk in the **My Computer** window. The status bar at the bottom of the window shows a disk capacity of 1.46 GB (.04 GB is used by the file system for internal configuration data) and 1.45 GB of free space (Figure 24).
12. For a mounted virtual disk, open Windows Explorer and access the `v:\` drive. The new virtual disk is displayed as a mounted drive (`v:\HQ Data`).

Refer to the SVR online help for more information.

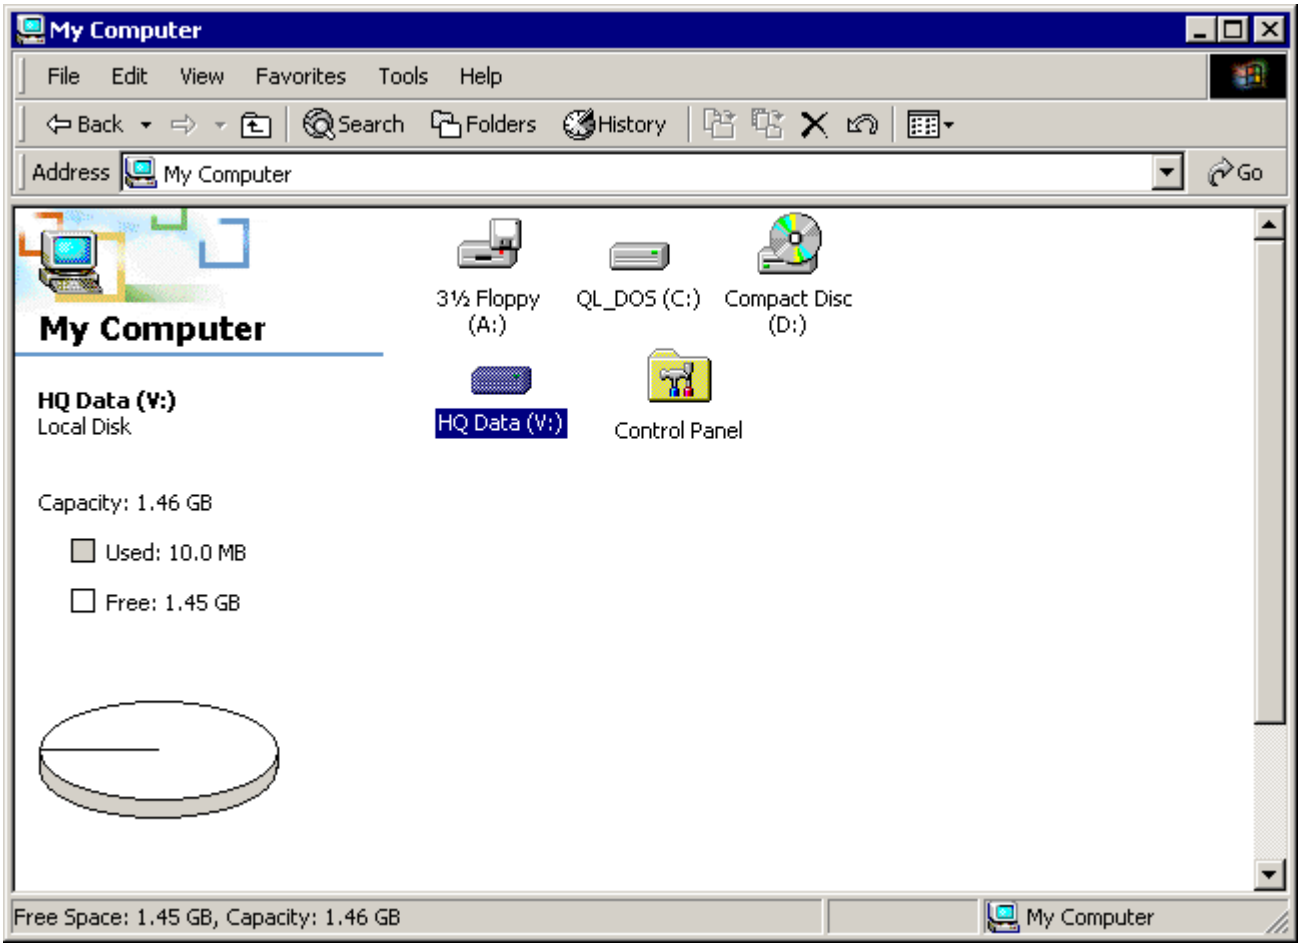

Figure 24: Viewing the virtual disk from My Computer

Mapping drive letters

During system startup, drive letters are automatically allocated from the beginning of the alphabet. The pool does not come online until later in the startup sequence. Therefore, if you selected a drive letter for a pool near the beginning of the alphabet, the drive letter may be allocated to another disk and will not be available for the virtual 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.

Note: Use SVR, not Disk Management, to map or unmap drive letters to virtual disks.

Formatting virtual disks

Note: You must map a drive letter to a virtual disk before you format it.

When you format a virtual disk, you can change the following NTFS disk parameters:

- **Allocation unit size**—The smallest amount of disk space that can be allocated to a file. SVR sets the default size to 4 KB to minimize the amount of lost space and the amount of fragmentation on the volume. HP recommends that you use the default size.
- **Volume label**—The name Windows uses to identify the virtual disk. If you change the volume label, it does not change the name of the virtual disk in SVR.

If you use Disk Management to format a virtual disk, ensure that you choose the NTFS file system. SVR does not support the FAT file system.

Caution: Do not use Disk Management to repartition a virtual disk; it may cause data loss.

See the SVR online help for more information.

Defragmenting virtual disks

When you defragment a virtual disk, you consolidate the space used by the files on the disk. This results in data being moved around on the disk.

If you defragment a virtual disk that contains a snapshot, moving the data causes copy-out operations. Although the data does not change, the contents of sectors on the disk change because the data is moving to new locations on the disk, resulting in unnecessary copy-out operations that degrade performance and waste disk space. To avoid this, delete all snapshots on a virtual disk before defragmenting. See “Snapshots overview” on page 82 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, SVR uses 4 KB as the default allocation unit size. See “Formatting virtual disks” on page 69 for more information.

Viewing virtual disk properties

You can view information about virtual disks using Replication Manager or My Computer in Windows. This section describes the following topics:

- Using Replication Manager
- Using My Computer

Using Replication Manager

To view information about a virtual disk in Replication Manager, select a pool in the scope pane. The results pane shows the following information:

- **Name**—The name of the virtual disk. The drive letter that is mapped to the virtual disk is shown in brackets.
- **Type**—The value is `Virtual Disk`.
- **Capacity**—The capacity of the virtual disk.
- **Delspace**—The amount of space you gain in the pool if you delete the virtual disk. The Delspace value is equal to the capacity of the virtual disk. If the virtual disk contains snapshots, the value is zero.You cannot delete a virtual disk if it contains snapshots.
- **Created**—The date on which the virtual disk was created.
- **Family**—The family to which the virtual disk belongs.
- **Owner**—The name of the cluster node that owns the pool resource. This information does not display for a stand–alone computer.

To view additional information about a virtual disk, right-click the virtual disk and select **Properties**. The Properties window opens (Figure 25).

The Properties window has three tabs:

- **General**—Displays the same information as Replication Manager. It shows the pool to which the virtual disk belongs, and the disk number and drive letter of the virtual disk.
- **Scheduled Tasks**—Lists any tasks that have been scheduled for the virtual disk. See "Scheduling wizards" on page 121 for more information.
- **Disk Map**—Shows how the virtual disk maps to the physical disk (Figure 26).

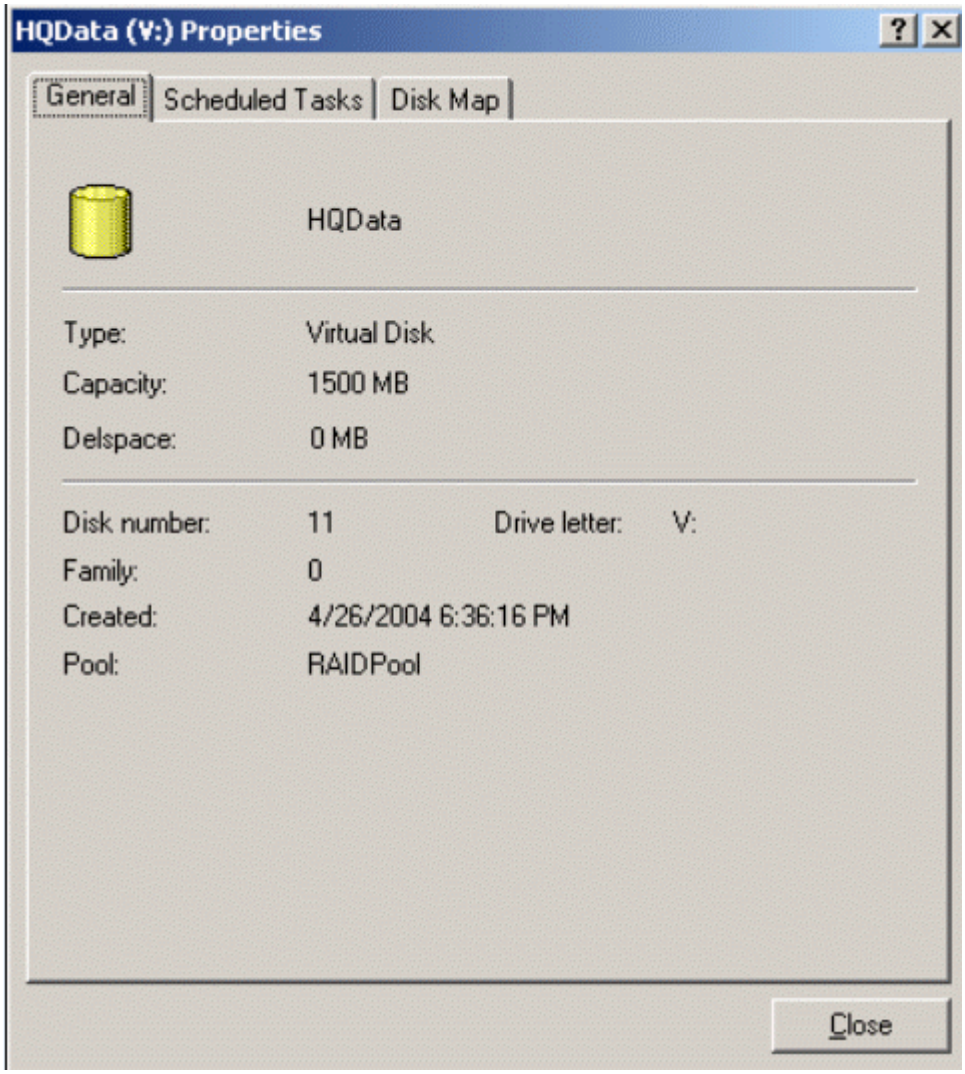

Figure 25: Properties window

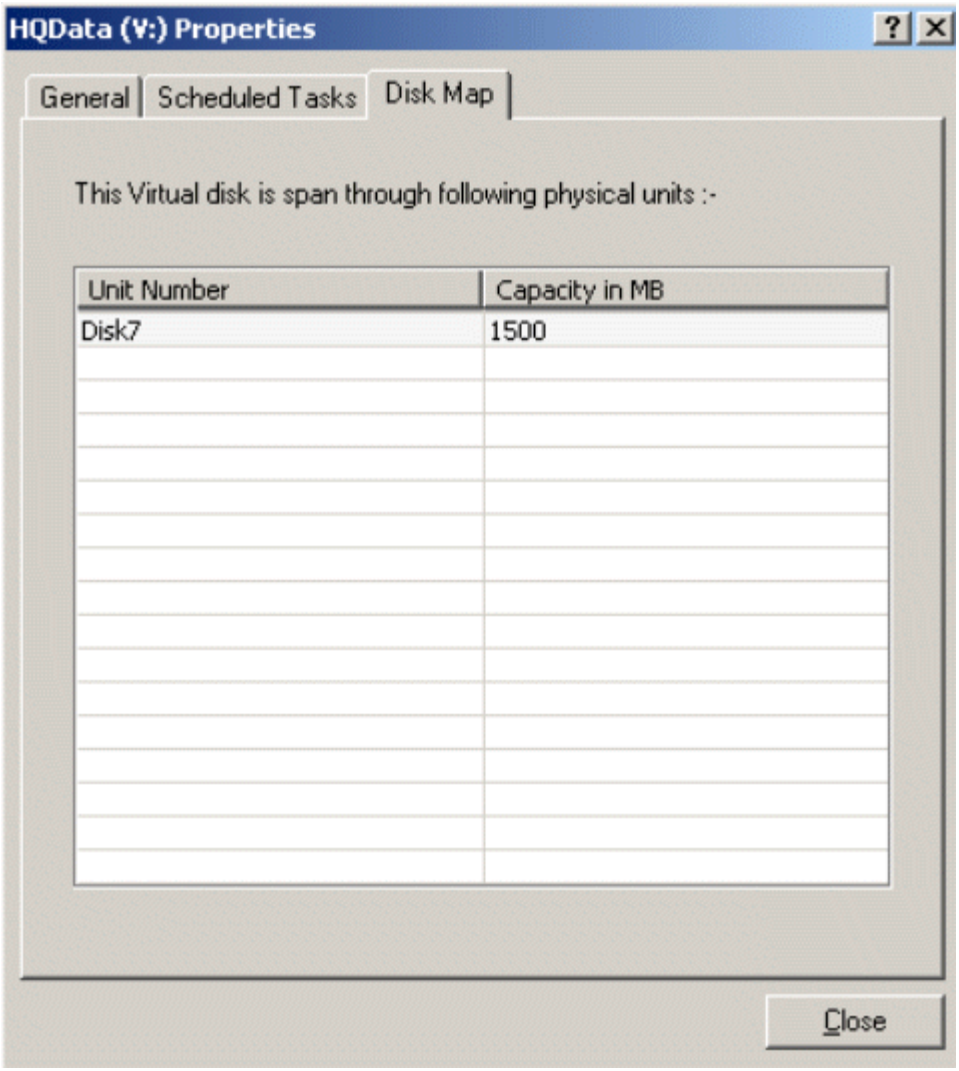

Figure 26: Disk Map window

Using My Computer

Virtual disks are displayed as icons in the **My Computer** window. Right–click a virtual disk and select **Properties** to display the Properties window for the virtual disk (Figure 27).

The General tab displays the following information:

- **Label**—The name Windows uses to identify the virtual disk. If you change the label, it does not change the name of the virtual disk in SVR.
- **Type**—The type of disk.
- **File system**—The file system for this volume.
- **Used space**—The amount of virtual disk space currently being used.
- **Free space**—The amount of virtual disk space 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 information as the SVR Properties window (Figure 25).

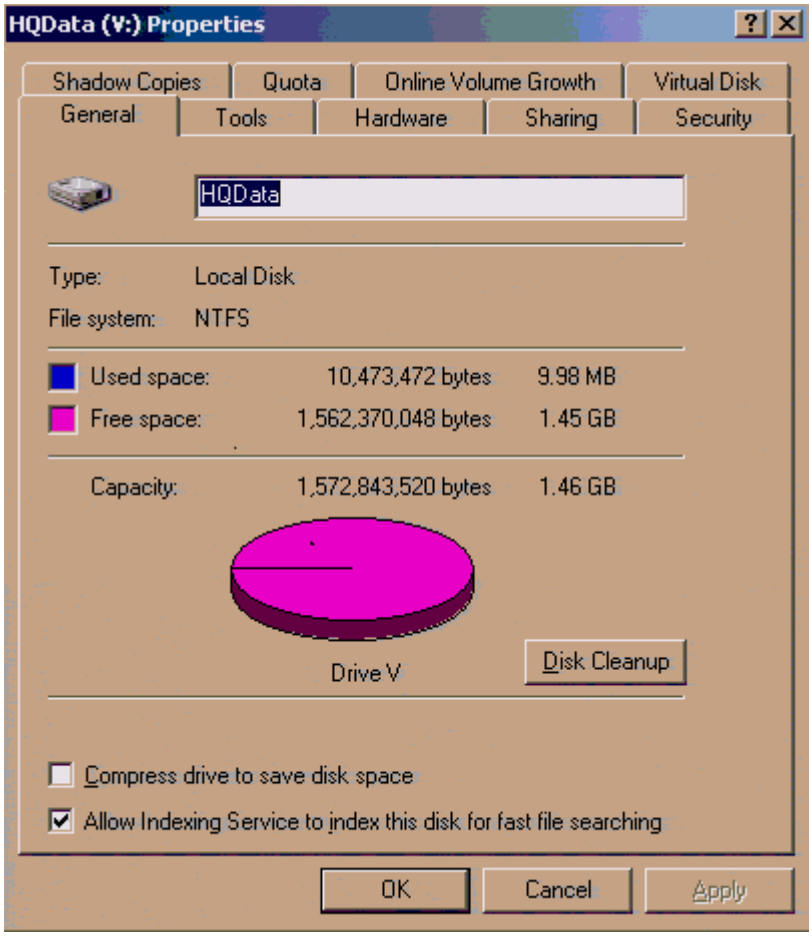

Figure 27: Properties window (My Computer)

Restoring virtual disks

When data on a virtual disk is corrupted, you can restore the virtual disk from a tape backup. You can also use SVR to restore virtual disks. This section describes the following topics:

- Using SnapBack
- Restoring virtual disks manually

Using SnapBack

You can use SnapBack to restore virtual disks automatically. SnapBack deletes all data from the virtual disk and replaces it with the snapshot's data. Using a snapshot to recover data prevents copying a file to a temporary location and consuming disk space.

Note: If data on the virtual disk changes between the time you made the snapshot and the time you perform SnapBack, the new data will be lost.

SnapBack performs the following tasks:

- Maps and remaps drives.
- Copies files.
- Deletes old snapshots or virtual disks.

SnapBack retains the source snapshot but deletes any snapshot that is older than the source snapshot. This ensures:

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

The duration of the SnapBack process may vary, depending on factors such as the amount of data on the snapshot, the number of storage units in the pool, and your hardware configuration. See "SnapBack scenario" on page 76 for more information.

SnapBack scenario

In this scenario, a virtual disk is corrupted during software testing. Table 4 shows the result of using SnapBack in a hardware and software environment indicated.

Table 4: Using Snapback

Using SnapBack	
Configuration	HP StorageWorks ProLiant server RA4200 SCSI controller External storage array with fourteen 18.2 GB drives 512 MB memory Two 500 MHz processors
SVR setting	Pool: 242 GB capacity One virtual disk (`vdisk_1`): 122 GB total capacity; 60 GB used One snapshot (`snap_1`) of the virtual disk
Event	Install new software for testing purposes on `vdisk_1`, write an additional 60 GB of data. (`snap_1's` Delspace increases to 60 GB.) The software corrupts the data on `vdisk_1`.
Response	Start SVR, 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**: SVR does not restore the 60 GB of data that was written to `vdisk_1` after the snapshot was created. That data is lost.

Restoring virtual disks manually

SnapBack is the most direct way to restore a virtual disk. However, if you want to perform a manual restore, use one of the following methods:

- **Method 1**—Unmap the drive letter of the virtual disk and map that drive letter to an existing snapshot of the virtual disk. This is the quickest method of the three, but it makes the original virtual disk an orphan. It also requires more pool space and reduces performance because of additional copy-outs.
- **Method 2**—Assign a drive letter to an existing snapshot and copy files from the snapshot to the virtual disk. Create a new snapshot of the virtual disk, then delete the old snapshot. This method is best for restoring files that have been deleted or corrupted.
- **Method 3**—Create a new virtual disk and copy files from an existing snapshot to the new virtual disk. When the copy is complete, delete the old virtual disk and all of its snapshots. This method is the manual equivalent of SnapBack and can be used to restore an entire volume. However, this method consumes more space temporarily because there are two virtual disk requiring pool space. To conserve disk and pool space, use SnapBack.

Deleting virtual disks

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

When deleting virtual disks, consider the following:

- Deleting a virtual disk deletes all data on the disk. If the virtual disk contains data you want to keep, back up the data first.
- You cannot delete a virtual disk that contains snapshots.
- After you delete a virtual disk, be sure to delete any scheduled tasks for the disk. If you reuse a virtual disk name, the scheduled tasks will run on the new virtual disk.

Refer to the SVR online help for more information.

Growing virtual disks

SVR provides an Online Volume Growth (OVG) wizard that lets you increase the capacity of virtual disks without having to restart the computer. See “Online Volume Growth” on page 105 for more information.

Note: You cannot use OVG to grow virtual disks that have associated snapshots. Before growing a virtual disk, you must delete all of its snapshots.

Managing Snapshots

6

This chapter describes how to create and manage snapshots:

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

Snapshots overview

A snapshot is a disk in a pool that you create 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 snapshots of those snapshots are known as a *family*. You can have up to 12 snapshots in each family.

Creating a snapshot is instantaneous because you do not actually copy any data. When you first create the snapshot, it holds exactly *the same* data as its parent disk, so there is no need to make a physical copy of the data. The snapshot uses disk space if you write data to the original virtual disk or the snapshot itself. See "Using disk space for snapshots" for more information.

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

Snapshots are useful when you need a quick copy of your production data without disrupting applications. You can use snapshots to do online backup and restore, test new applications, or populate your data warehouse or web server.

Volume Shadow copy Service

Volume Shadow copy Service (VSS) is a feature of Microsoft Windows Server 2003 that provides a general infrastructure for creating point-in-time copies of data on a volume. You can use SVR to create VSS snapshots on a Windows system. Refer to the Microsoft Volume Shadow copy Service documentation for more information.

Using disk space for snapshots

Initially, a snapshot shares the same physical disk space as its parent. If you create, modify, or delete files or directories stored on the snapshot or parent disk, SVR must first copy the existing data before changes are made. This operation is called a *copy-out*.

When a copy-out occurs, the snapshot uses disk space in units called *segments*. A segment is a fixed-size unit of contiguous bytes of disk space. You set the segment size when you create a pool (see "Managing Pools" on page 41). Each copy-out operation uses a segment of disk space in the pool.

For example, you create a snapshot that contains a file called `small.dat` It is identical to the `small.dat` file on the parent disk. Both files share the same disk space (Figure 28).

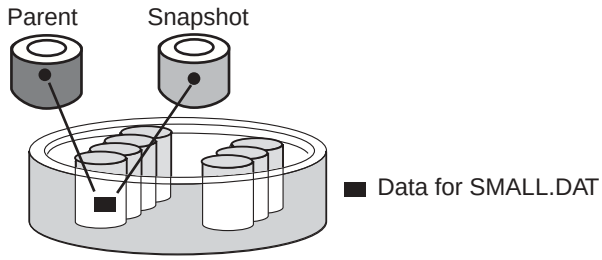

Figure 28: Example of a new snapshot using disk space

You then write data to the parent disk's `small.dat` file. Before the data is changed, SVR creates a separate physical copy of the old data (copy-out operation). SVR then updates the parent disk with the new data (Figure 29).

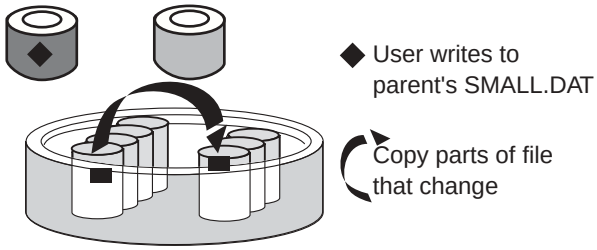

Figure 29: Copy-out operation

The snapshot's `small.dat` file contains the old data and the parent disk's `small.dat` file contains the new data (Figure 30). Because the `small.dat` file is small, it can be copied as one 128 KB segment. For larger files, the modified data is copied in multiple segments.

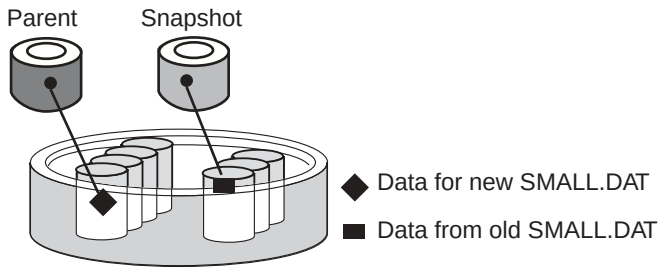

Figure 30: Disk space usage after copy-out

If the same segment of the `small.dat` file is changed again, SVR does not have to create another copy. The snapshot has a copy of the original segment, which does not change; therefore, there is no need to copy it again.

Creating snapshots

When creating snapshots, consider the following:

- You can only create snapshots after you create a virtual disk.
- Do not create snapshots when there is heavy I/O activity on the virtual disk. Doing so can increase the time needed to create the snapshot.
- When using advanced applications, such as database products, you must assess the application's requirements for data consistency. Before you create snapshots, you may need to pause the application to ensure data consistency.

To create a snapshot:

1. In the scope pane, right–click a virtual disk and select **New** > **Snapshot**.

   The New Snapshot Wizard window opens.

2. Click **Next**.

   The Snapshot Information window opens (Figure 31).

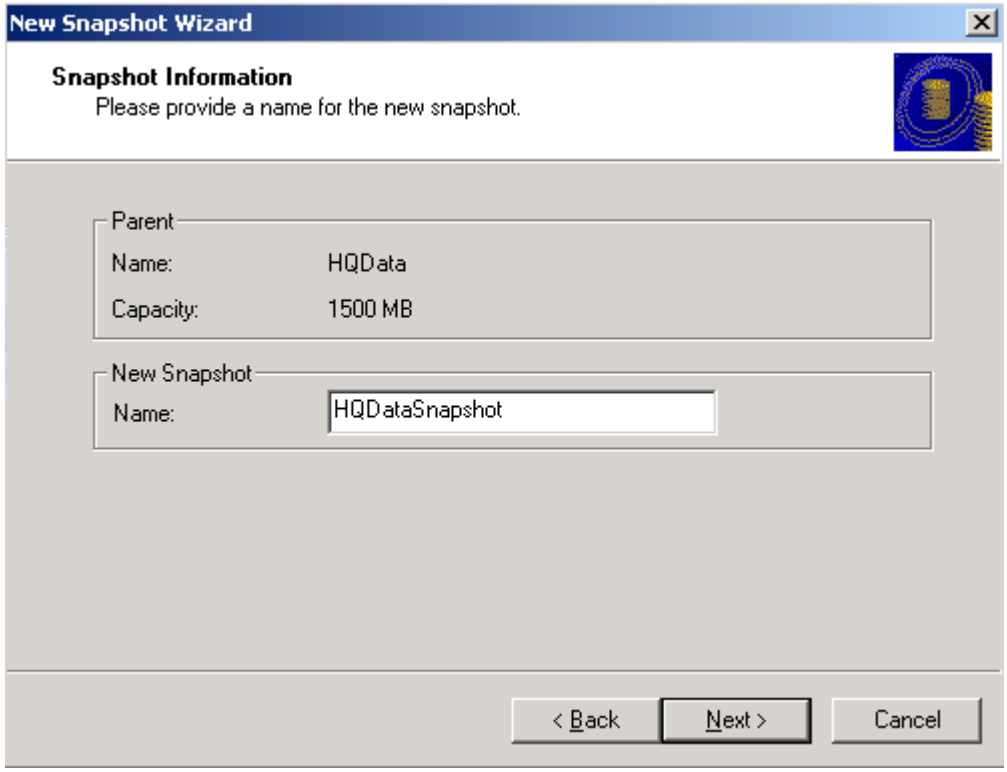

Figure 31: Snapshot Information window

3. Enter a name for the snapshot in the Name box and click **Next**.

   The Drive Letter Assignment window opens.

4. Do one of the following:

   - Select **Map drive letter** and choose the drive letter you want to map to the snapshot (Figure 32). Mapping drive letters is the same for snapshots and virtual disks. See "Mapping drive letters" on page 68 for more information.

     By default, the volume label for the snapshot is the same as that of the parent disk. To differentiate this snapshot from other snapshots and the parent virtual disk, edit the name in the Volume Label box.

   - Click **Mount this volume at an empty folder**. Click **Browse** and choose an empty folder on an NTFS volume to mount the drive (Figure 33). Click **OK**.

   - Select **No drive letter**. The snapshot remains offline until you map a drive letter or create a mount point to it.

5. To create a VSS snapshot, select **Enable Volume Shadow Copy**. This option only displays if you are using Windows Server 2003.

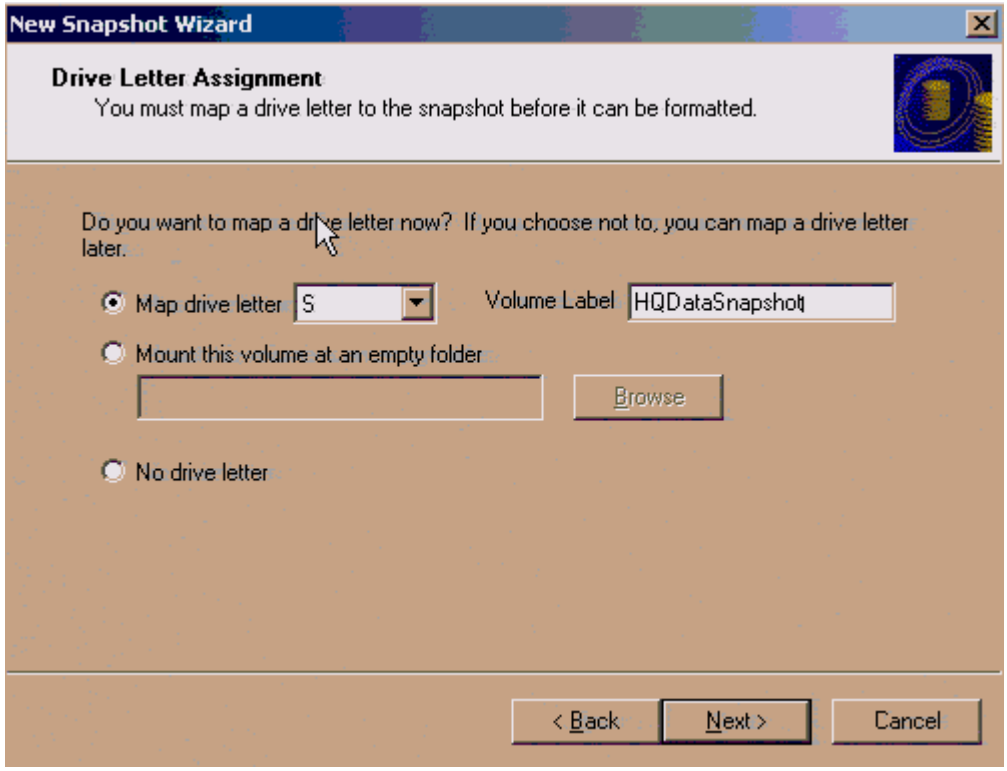

Figure 32: Mapping a drive letter

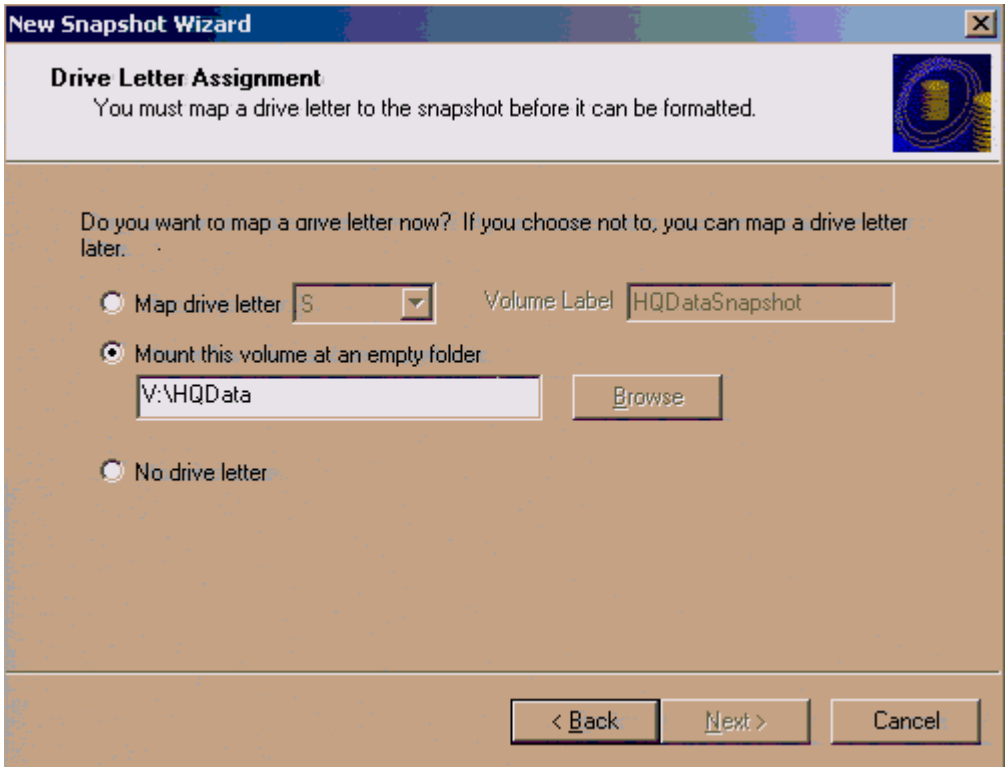

Figure 33: Mounting a volume

6. Click **Next**.

   The Completing the New Snapshot Wizard window opens.

7. Verify that there are no error messages and click **Finish**.

8. Open Replication Manager.

   The new snapshot is listed under its parent virtual disk in the scope pane (Figure 34).

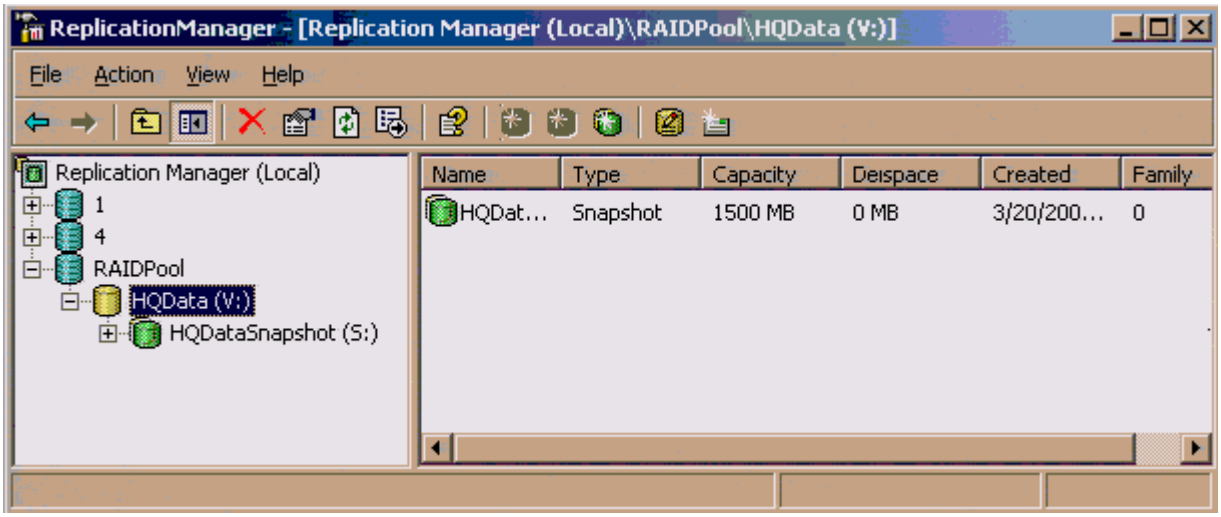

Figure 34: Viewing the new snapshot

9. Open My Computer to view the new snapshot.
10. Click the snapshot to view its free space and capacity in the bottom status bar (Figure 35). It has the same free space as its parent disk because it contains the same files.

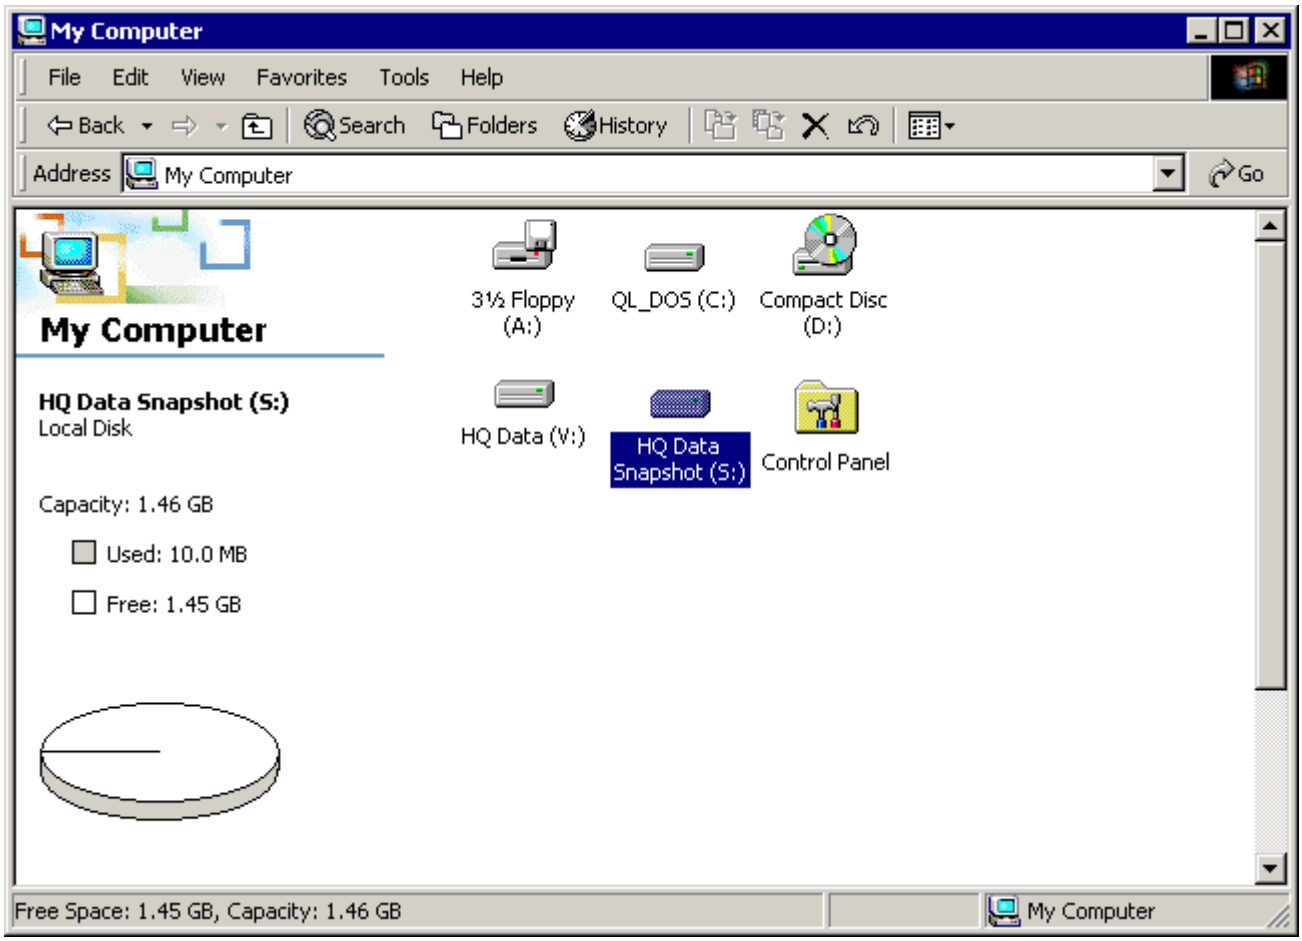

Figure 35: Viewing the snapshot from My Computer

11. Delete the `myfile.txt` file on the parent disk (Figure 36).

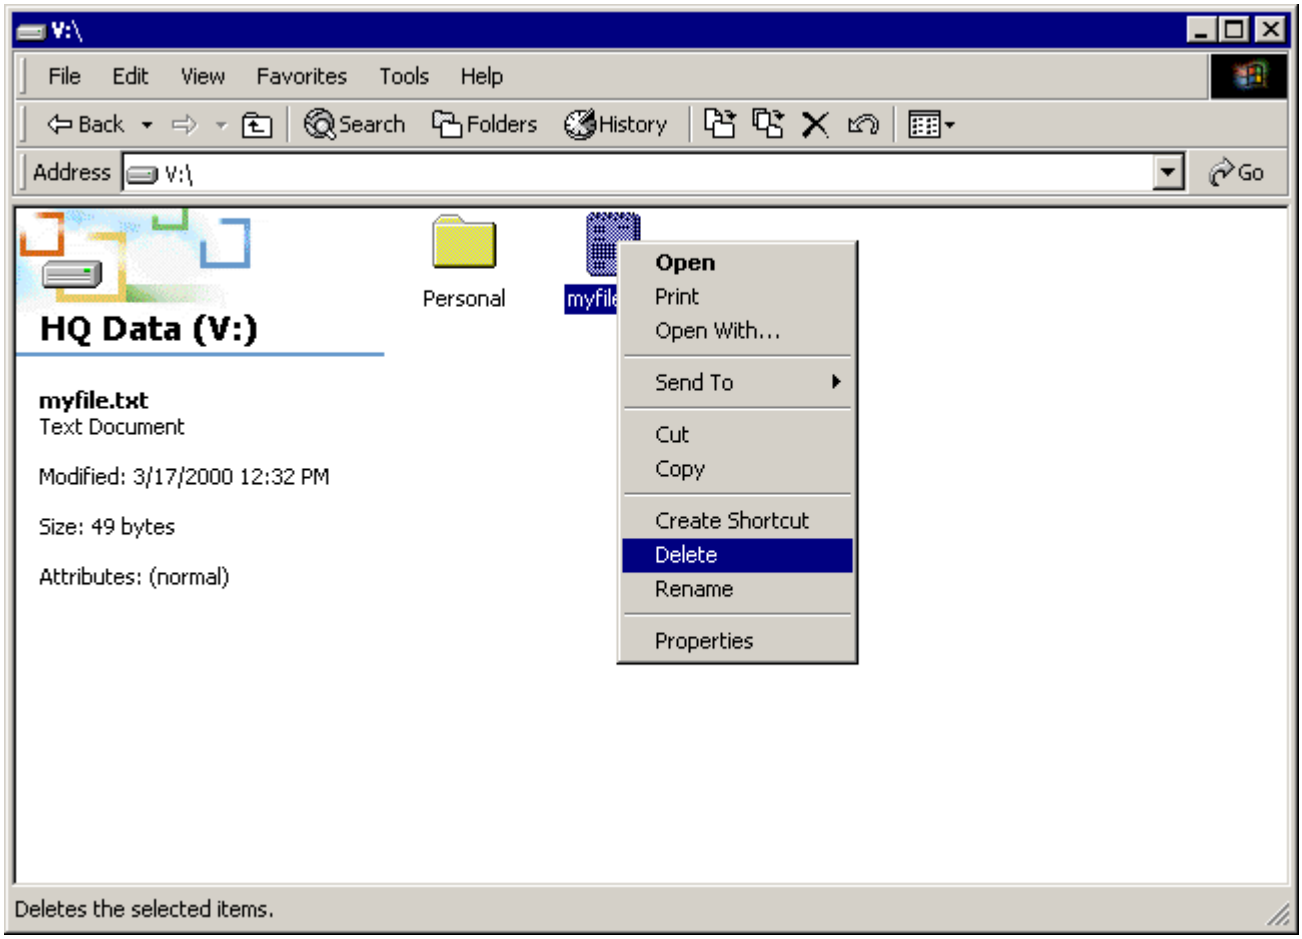

Figure 36: Deleting the myfile.txt file on the parent disk

The `myfile.txt` file still exists on the snapshot (Figure 37).

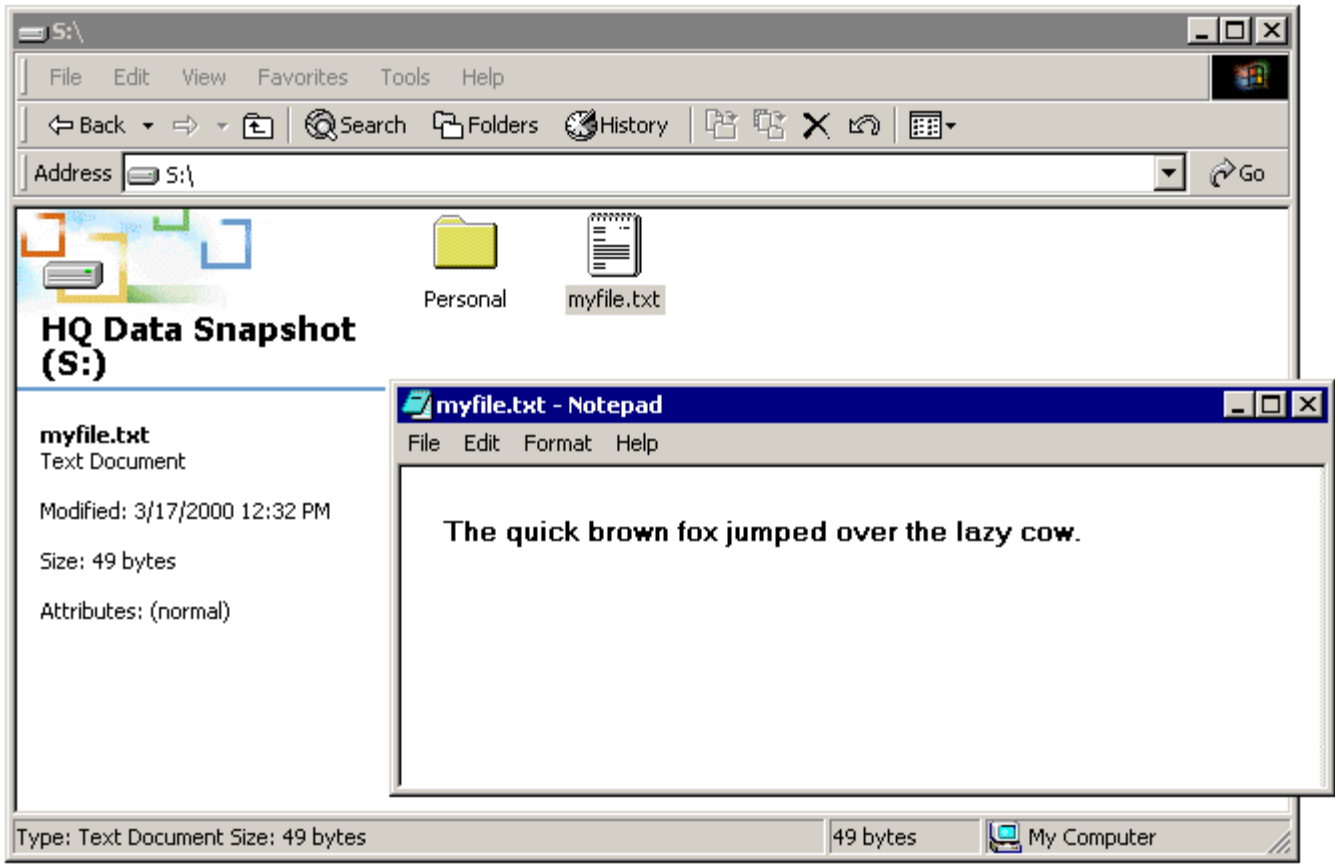

Figure 37: Viewing the myfile.txt file from the snapshot

12. Open Replication Manager and verify the pool's free space. The free space remains the same because the snapshot is not using disk space (Figure 38).The only update that occurred on the disk was an update to the file system tables, which uses very little space.

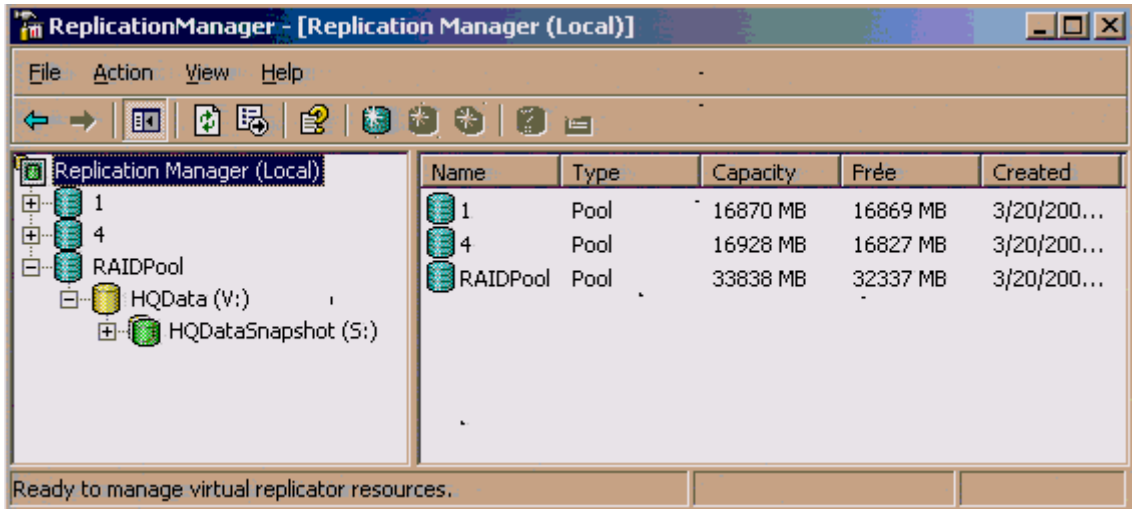

Figure 38: Verifying the pool's free space

Best practices

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

When data on a parent disk changes, segments are copied out to each of the parent disk's snapshots. SVR tracks copy-outs by associating an address with each segment of data. During startup, SVR pools "rebind" by reestablishing these address connections. The rebinding process is usually quick, but can take longer if the snapshots have large amounts of changed data (copy-outs).

You should not consider SVR snapshots your only means of backup. HP strongly recommends that you use SVR in conjunction with a standard backup tool, and that you follow these guidelines:

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

See "Using snapshots for backup and restore" on page 97 for more information on the best practices for using snapshots to back up virtual disks. See "Exceptions" on page 26 for additional guidelines.

Using VSS snapshots

When creating VSS-assisted snapshots, consider the following:

- Each VSS-assisted snapshots takes approximately one and a half minutes to create during normal I/O activity on the parent disk.
- Do not create VSS-assisted snapshots during heavy I/O activity on the parent disk.
- Applications can register themselves as writers with the VSS infrastructure and can participate in the VSS snapshot creation. Before creating a VSS-assisted snapshot, ensure that the writer status is Stable; if not, the snapshot creation may fail.

  To check the writer status, open a command prompt window and enter:

  `"vssadmin list writers"`
- Do not create VSS-assisted snapshots with SVR and Windows Shadow Copies simultaneously. The snapshot or the Shadow Copy will fail.

Snapshots in a cluster

SVR does not create cluster resources for snapshots as it does for pools. The snapshots in a pool automatically appear on the node that owns the pool resource.

Scheduling snapshots

You can the SVR Task Scheduling wizards to schedule snapshots of virtual disks or other snapshots. See “Scheduling Tasks” on page 119 for more information.

You can also use the `SmartSnap` utility to specify the number of snapshots maintained for a virtual disk. Refer to the *HP OpenView Storage Virtual Replicator Command Line Interface Reference Guide* for more information.

Viewing snapshot information

You can use Replication Manager to view information for snapshots. When you select a virtual disk in the scope pane, the results pane displays the following information for each snapshot in the family:

- **Name**—The name of the snapshot. The drive letter mapped to the snapshot is shown in brackets.
- **Type**—The value is `Snapshot`.
- **Capacity**—The capacity of the snapshot. The file system uses this value. For example, Windows Explorer uses this value to display the snapshot's size. Capacity also represents the maximum pool space the snapshot can use if you modified the entire contents of the snapshot or the parent disk.
- **Delspace**—The amount of space you will gain in the pool if you delete the snapshot. The value of Delspace is the amount of pool space used exclusively by this snapshot, not space that is shared with other snapshots in the family.
- **Created**—The date on which the snapshot was created.
- **Family**—The family to which the snapshot belongs.
- **Owner**—The name of the cluster node that owns the pool resource. This information does not display for a stand–alone computer.

To view additional information about a snapshot, right–click the snapshot and select **Properties**. The Properties window opens (Figure 39).

The Properties window contains two tabs:

- **General**—Displays the same information as Replication Manager. It shows the pool to which the snapshot belongs and the disk number and drive letter of the snapshot. The disk number may change when you restart the computer or when the pool resource fails over within the cluster.
- **Scheduled Tasks**—Lists tasks that have been scheduled for the snapshot. See "Scheduling wizards" on page 121 for more information.

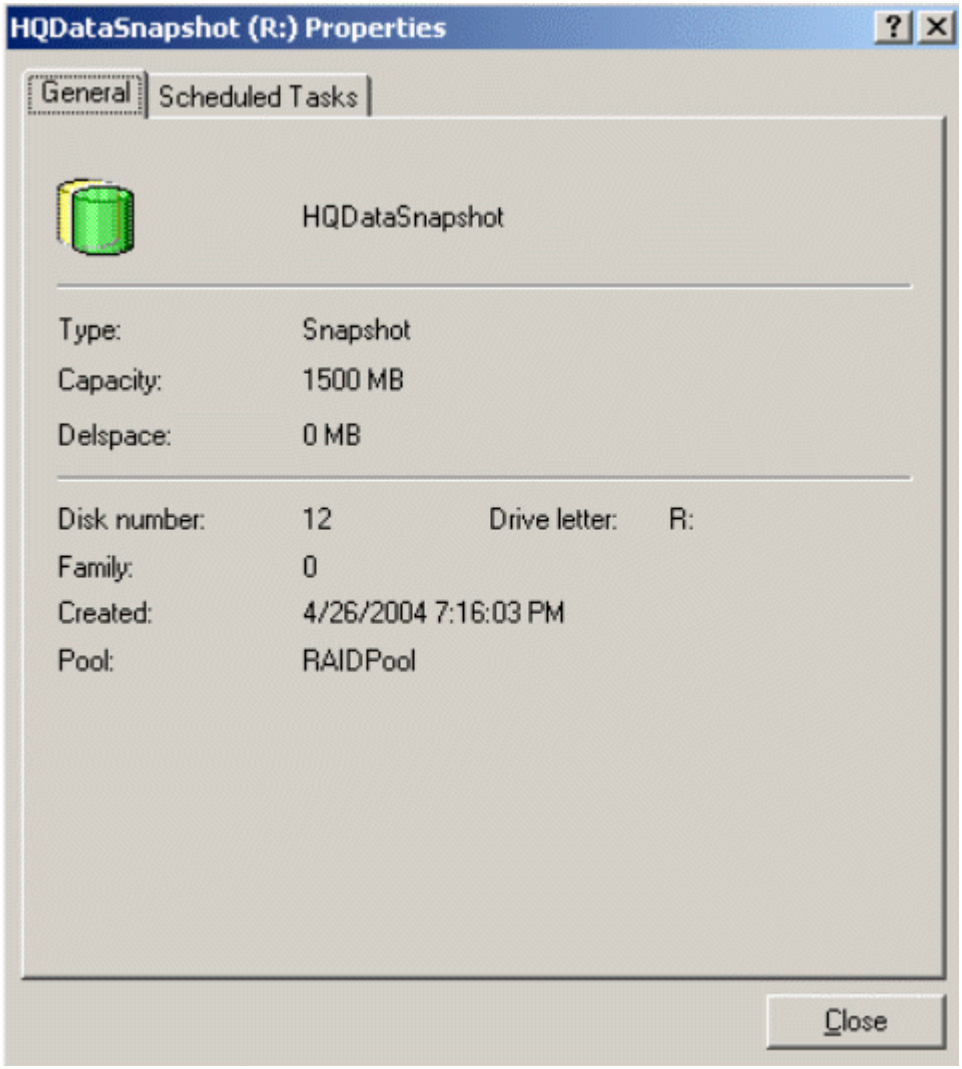

Figure 39: Properties window

Using snapshots for backup and restore

To use a snapshot to back up a virtual disk:

1. Stop the application.
2. Create a snapshot.
3. Restart the application.
4. Back up the virtual disk.
5. Delete the snapshot created in step 2.

To automate this process, you can:

- Use `SnapMgr` to create batch jobs and automate backups.
- Use the Snapshot for Backup wizard. The wizard sends an e-mail or pager notification when the backup is complete.

Refer to the SVR online help for more information.

Using snapshots to restore virtual disks

You can use snapshots of virtual disks to restore individual files or the entire volume if the disk is corrupted. See “Restoring virtual disks” on page 75 for more information.

Other backup and restore solutions

The HP storage web site contains information about using snapshots for backups and restores in specific applications, such as Microsoft Exchange. For more information, go to http://h18000.www1.hp.com/products/sanworks/vr/index.html.

Deleting snapshots

You can delete a snapshot that you no longer need to gain space in the pool. The pool's free space increases by the snapshot's Delspace.

When deleting snapshots, consider the following:

- If the snapshot shares disk space with another snapshot, the Delspace of the other snapshot increases by the amount of this disk space.
- If the snapshot is the only snapshot in the family, the Delspace of the family's virtual disk increases from zero to the capacity of the virtual disk. If you delete the virtual disk, its Delspace becomes the same as its capacity.
- Deleting a snapshot deletes all data stored on the snapshot, but does not affect the data stored on its parent disk or any other disks in the family.
- After you delete a snapshot, delete any scheduled tasks for the snapshot. Otherwise, if you reuse the snapshot name, the old tasks will run on the new snapshot.

See “Scheduling Tasks” on page 119 and “Managing SVR Policies” on page 135 for more information about deleting snapshots.

Migrating Data

7

This chapter describes how to use the data migration wizard:

Overview

The Data Migration wizard enables you to move data from one physical disk to another in a pool without affecting the virtual storage environment. For example, you can migrate data to new, higher-performance disks without deleting the pool or virtual disks and while maintaining I/O access to the pool. When data migration is complete, you can remove the old disks. You can also migrate data from multiple small disks to one large disk.

Caution: HP strongly recommends that you plan data migration properly. The process may take a long time to complete, depending on the size of the data being migrated. Use the Disk Map feature to determine the size of the virtual disk residing on the source disk being migrated.

You can pause, resume, or cancel the data migration at any time. These actions are recorded in the event log. A window displays during the operation so you can monitor the progress of the migration.

Guidelines

When migrating data, consider the following:

- The capacity of the destination disk should be greater or equal to the combined capacity of the source disk's data being migrated.
- You must delete all snapshots on the pool. Any scheduled snapshots will fail during data migration.
- Although I/O activity to the virtual disks is allowed, you cannot perform any management tasks on the pool during data migration. You must pause the data migration to complete these tasks. When you have finished, you can resume data migration. You can also cancel the migration, which would leave the virtual disk spanning the source and destination disks.
- Data migration is designed for disaster recovery and for failovers in a clustered environment. If a system failure or failover occurs, the data migration process is not automatically resumed after system recovery. You must restart the migration to complete it. System recovery only ensures that the pool is not corrupted and is in a stable state.
- Data migration is limited to physical disks in the pool. You cannot start another migration within the same pool until the first migration is completed.

Using the Data Migration wizard

To migrate data to another pool:

1. In the scope pane, right–click the pool and select **All Tasks > Migrate Data**.

   The welcome window for the Data Migration wizard opens.

2. Click **Next**.

   The Data Migration window opens (Figure 40).

3. Select the type of target disk you want to use under **Target list should be**:

   - **Only Raw disks**—Add the disk to the pool before starting data migration.
   - **Only Pool disks**—Select from existing disks in the pool.
   - **Both Type**—Select from both raw and pool disks.

   Depending on your selection, the available disks are moved from the left column to the right column.

4. In the right column, select the target disk of the data migration.

5. If you want to delete the source disk after data migration is complete, select **Remove source disk after migration**.

6. Select a size under **Select Chunk size** or accept the default.

   The chunk size is the basic unit of migration. It is the number of blocks that can be migrated at one time.

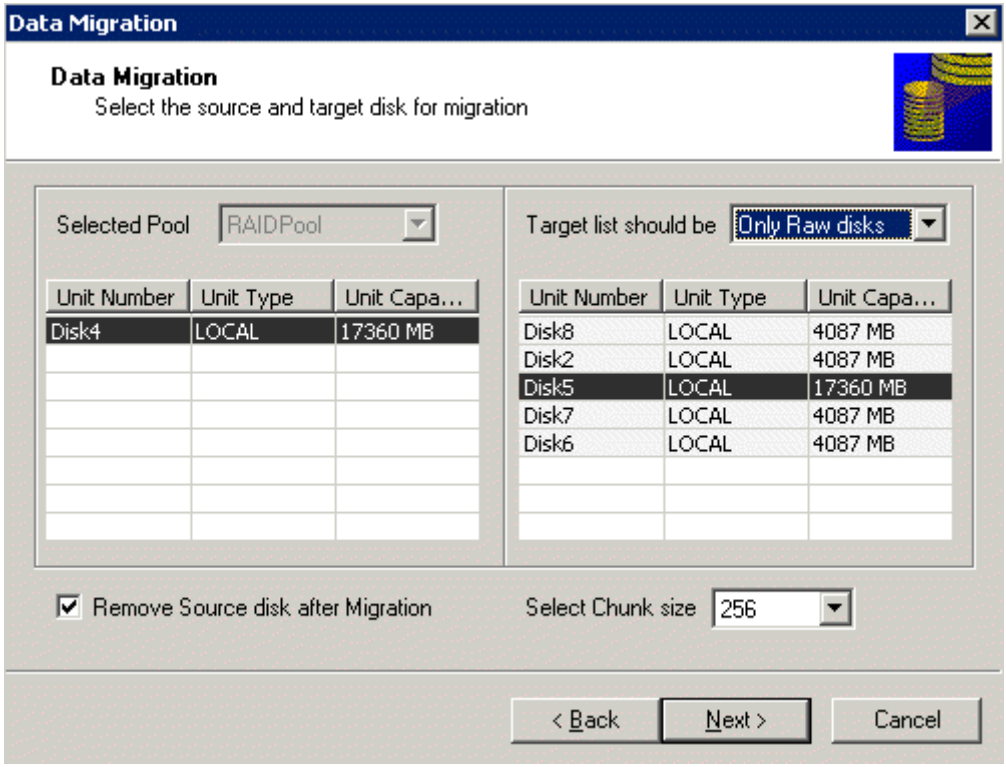

Figure 40: Data Migration window

7. Click **Next**.

   The completion window of the data migration window opens.

8. Click **Finish**.

   A status window opens to track the progress of the data migration (Figure 41).

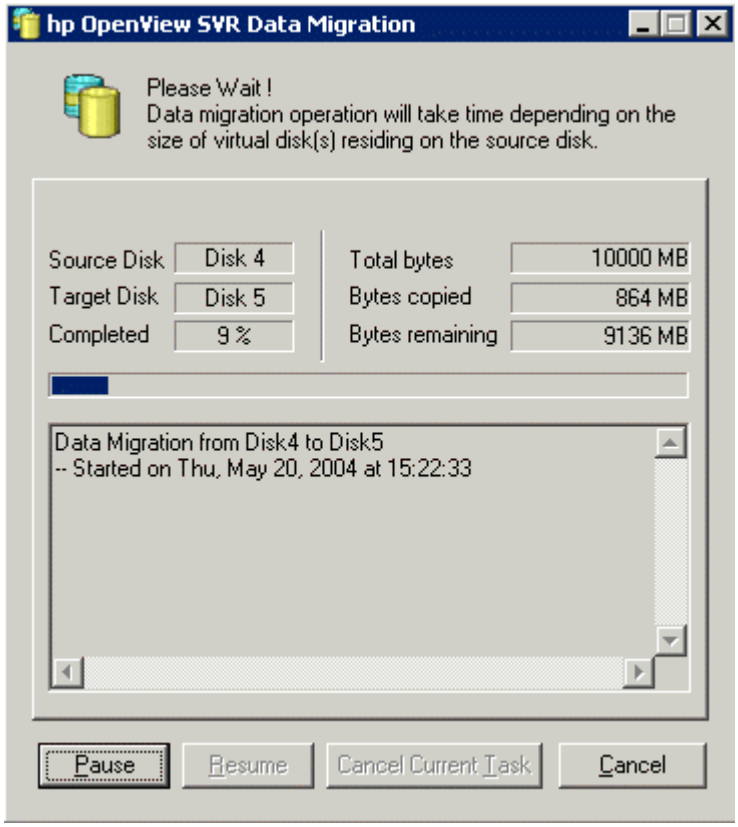

Figure 41: Status window

Online Volume Growth

8

This chapter provides information and examples to help you use Online Volume Growth (OVG). It describes the following topics:

Online Volume Growth overview

Using SVR's Online Volume Growth feature, you can increase storage capacity without restarting the computer. You can grow a volume on a basic (physical) disk by creating additional space on the disk. You can grow a volume on a virtual disk by adding storage units to the pool and then increasing the virtual disk's size.

Refer to the Online Volume Growth online help for detailed instructions.

Understanding disk types

You can use Online Volume Growth with basic and virtual disks formatted with NTFS. This section describes the differences between basic, virtual, and dynamic disks.

Note: Online Volume Growth does not support Windows dynamic disks.

Basic disks

A basic disk is a partitioned disk that contains a partition table and a maximum of four primary partitions. To create a configuration with more than four partitions, create an extended partition entry (reducing the possible number of primary partitions to three) and one or more logical drives.

You can format each primary partition or logical drive with a file system and use it as a volume in the operating system.

Table 5 describes the configuration for a basic disk.

Table 5: 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)

Simplicity is one benefit of using basic disks for Online Volume Growth. After performing multiple Online Volume Growth operations, the result is a single partition that has grown over time.

Table 6 describes the configuration for a basic disk after using Online Volume Growth.

Table 6: Basic disk configuration after Online Volume Growth

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

Several operating systems recognize partitions, including MS-DOS, Windows 3.1, Windows 95/98, and Windows NT. Use basic disks to ensure backwards compatibility. For example, if you dual-boot a system between Windows 95/98 and Windows 2000, you must use basic disks to support volumes in both environments.

Basic disks do not support volume configurations that span disk controllers, 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. Refer to the Microsoft Windows documentation for more information about basic disks and volumes.

Virtual disks

Using SVR, you can combine basic disks into pools of disk space and create virtual disks from those pools. Each virtual disk contains one partition formatted with NTFS.

Using Online Volume Growth, you can expand virtual disks without restarting the computer. For example, you can add physical disks to a RAID set, assign the new storage units to an existing pool, and then use the additional capacity to expand any virtual disks associated with the pool.

Dynamic disks

Windows supports dynamic disks, which do not require Online Volume Growth to increase in size. Instead of partition table, dynamic disks use a proprietary configuration database located at the end of the physical disk.

There are certain benefits unavailable with dynamic disks, such as extending a logical drive with an HP RAID controller. To extend a volume on a dynamic disk, you must create a new logical drive on an existing controller or use a new controller.

Although you can convert a basic disk to a dynamic disk and retain volumes on the disk, you cannot revert to a basic disk without first deleting the volumes on the disk. Therefore, when using Online Volume Growth, *do not convert eligible disks from basic to dynamic*.

Refer to the Microsoft Windows documentation for more information about dynamic disks and dynamic volumes.

Planning Online Volume Growth

When growing disks, consider the following:

- Make a backup copy of the data.
- Plan for new storage. Adding capacity may increase the time required to complete backups.
- Ensure that pools have adequate free space. Increasing the size of a virtual disk uses more space in the pool, which reduces the amount of free space available for snapshots.
- When you replace hard drives or use the HP Array Configuration Utility to expand arrays or extend logical drives, ensure that the task is complete before using Online Volume Growth.

  Using the HP Array Configuration Utility to expand an array or extend a logical drive takes time. Examine the priority settings for the controller. Refer to the HP Array Configuration Utility documentation for more information.
- Verify the volume configuration compatibility:
  - — You can only use Online Volume Growth with basic and virtual disks formatted with NTFS. Disks with NTFS volumes that are not configured for fault-tolerance are supported.
  - — Online Volume Growth is not compatible with dynamic disks, volume sets, mirror sets, stripe sets, or stripe sets with parity, or FAT and FAT32 volumes.

  You can convert FAT or FAT32 volumes to NTFS using the Windows `CONVERT` utility. Refer to the appropriate Windows documentation for more information.
- You must have Full Control permission at the top level of the volume you want to grow. By default, administrators have Full Control access to NTFS volumes in Windows. You do not need to modify this permission unless it has been explicitly denied.

Growing virtual disks

You can grow a virtual disk by:

- Selecting the disk in the Replication Manager MMC interface.
- Selecting the disk in Windows Explorer.
- Using the Online Volume Growth wizard.
- Using the command line interface.

Preparing for virtual disk growth

A pool must have 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. A pool can have up to 12 storage units. You can add storage units to a pool at any time, even when users are accessing its virtual disks and snapshots.

Once you add a storage unit to a pool, its size within the pool is fixed. Therefore, pool capacity does not increase when you extend a basic disk. You can only grow a pool by adding storage units to it.

After you add storage units to a pool, 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.

Refer to the the SVR online help for 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.

Growing basic disks

You can grow a basic disk by:

- Using the Online Volume Growth wizard.
- Selecting the disk in Windows Explorer.

Preparing for basic disk growth

Before you grow a volume on a basic disk, ensure that there is available space 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.

When growing basic disks, HP recommends that you set full access permission to one of the following groups:

- All users
- Administrators
- System

Table 7 shows a configuration that allows a **Windows 2000** volume to grow into the adjacent, unused space.

Table 7: Disk with 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

You can use Online Volume Growth to grow the `C:\` drive into the adjacent, unused space.

Table 8 shows a configuration without adjacent, unused space. The only way to increase space is to delete the **OLD TOOLS** volume.

Table 8: Disk without 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)	**OLD TOOLS (D:)** 10.00 GB FAT

Creating adjacent space on basic disks

You can create adjacent space on basic disks by:

- Deleting a partition
- Extending a logical drive

Deleting a partition

If you have a partition adjacent to the volume that you want to grow, you can delete this partition to increase space. Be sure to migrate any necessary data from the partition before deleting it. Table 9 shows that you grow the **Windows 2000** volume by deleting the **OLD TOOLS** volume and removing the partition.

Table 9: Deleting a 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:)]

Extending a logical drive

You can use the HP Array Configuration Utility to extend a logical drive on HP array controllers that support logical drive extension. 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.

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 (Table 10).

Table 10:

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

Table 10:

Initial configuration	
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/ID	0	1	2	3	4	5	6	7	8	9	10	11	12	13
0 (Internal)	4.3 GB	4.3 GB	4.3 GB	n/a	n/a	n/a			Bays: 6 internal/14 external					
1 (External)	4.3 GB	4.3 GB	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a

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
- Grow the boot volume using SVR

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.

Table 11 lists the configuration results after the upgrade.

Table 11:

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	n/a	n/a	n/a			Bays: 6 internal/14 external					
1 (External)	9.1 GB	4.3 GB	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a

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, select **Computer Management > Disk Management** to see the basic disk's new layout (Figure 42).

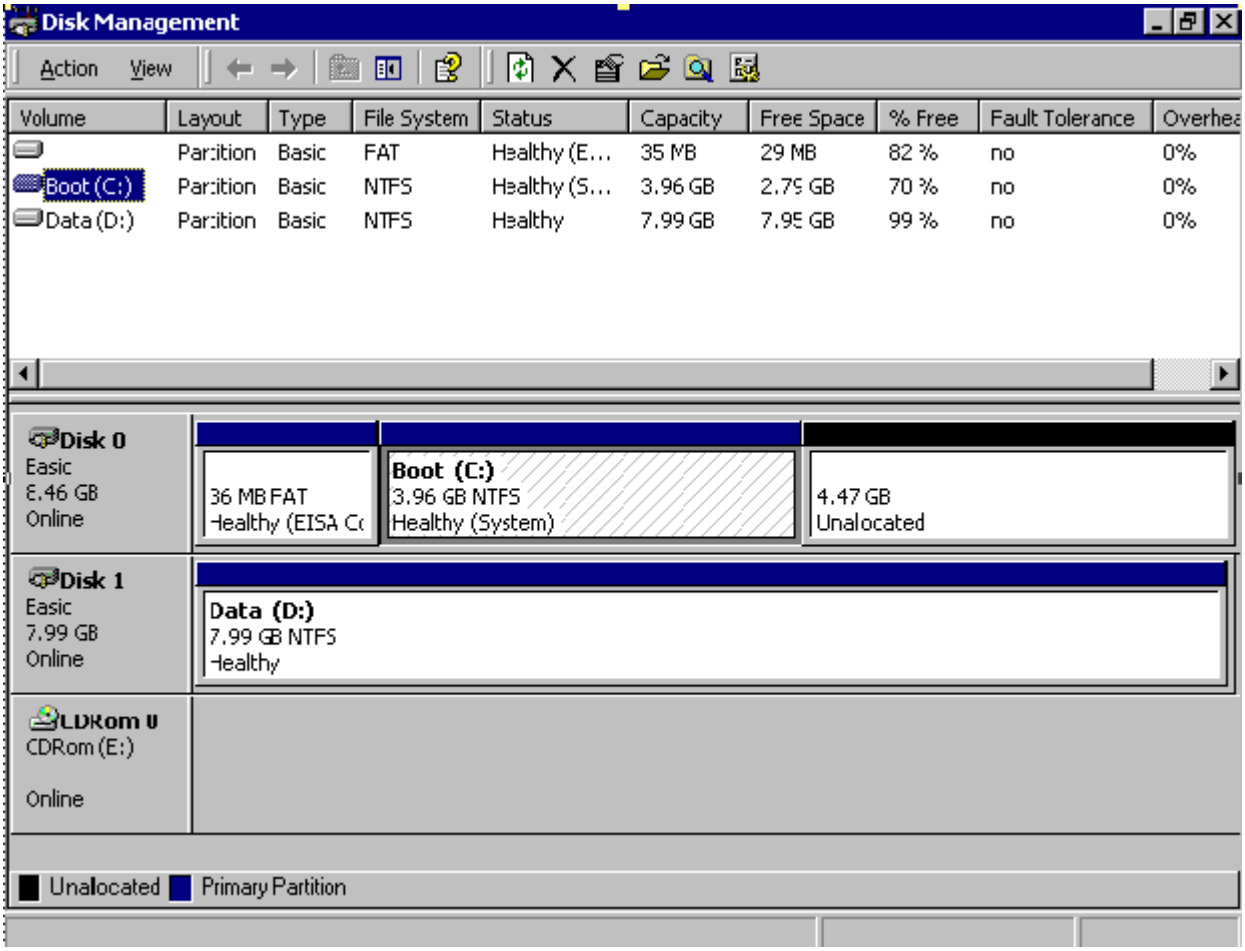

Figure 42: Disk Management window

Grow the boot volume using SVR

To grow the boot volume using SVR:

1. Select S**tart > All Programs > hp OpenView Storage Virtual Replicator > Online Volume Growth Wizard**.

   The Online Volume Growth Wizard window opens (Figure 43).

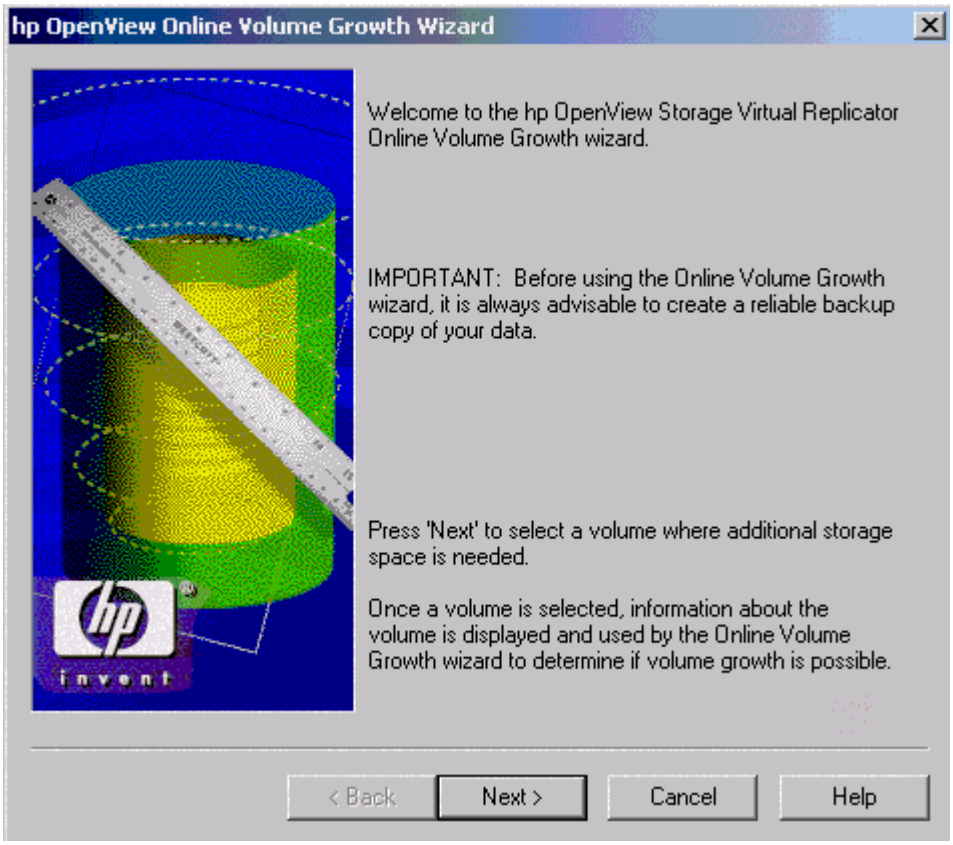

Figure 43: Online Volume Growth wizard

2. Click **Next** and follow the instructions in the wizard to grow the boot volume.
3. Select **Computer Management > Disk Management** to view the result. The boot volume has grown from 3.96 GB to 8.43 GB without causing any downtime (Figure 44).

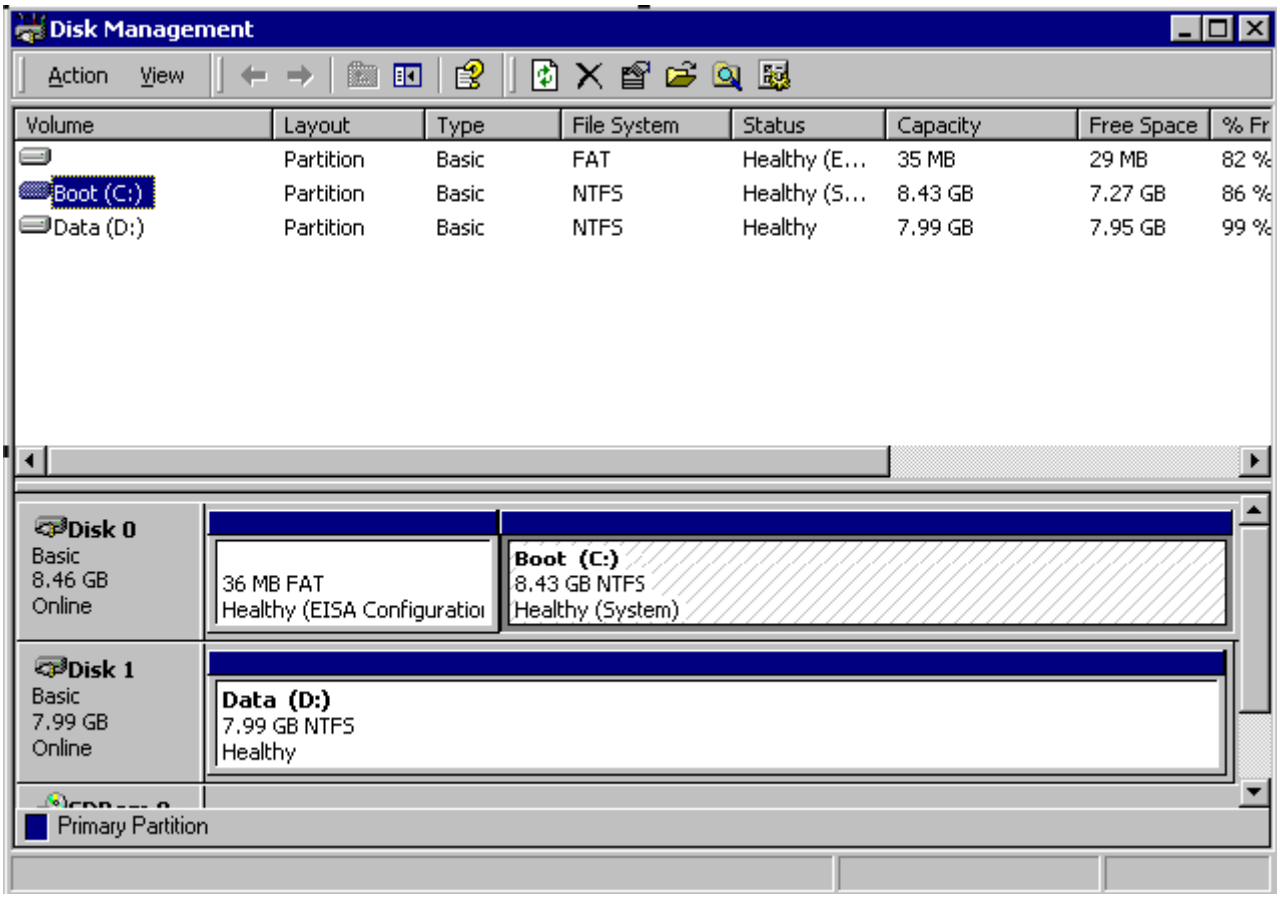

Figure 44: Disk Management window (grown volume)

Accessing Online Volume Growth from Windows Explorer

You can access the Online Volume Growth wizard from Windows Explorer by doing one of the following:

- Right–click the virtual or basic disk and select **Volume Growth**.
- Select the Online Volume Growth tab in the Properties window of the virtual or basic disk.

Refer to the Online Volume Growth Wizard online help for more information.

Scheduling Tasks

9

This chapter describes how to use the scheduling wizards to schedule tasks:

Scheduling wizards overview

SVR provides a set of wizards that automate virtual disk and snapshot tasks. You can access the scheduling wizards after you create a virtual disk.

SVR uses the Windows Scheduled Tasks applet to create and execute the schedules for the specified tasks. To use the scheduling wizards, ensure that the the Scheduled Tasks service is running and you have administrative rights.

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

Refer to the Scheduling Wizards online help for more information.

Scheduling wizards

This section describes the scheduling wizards.

Create snapshot

Use the Create Snapshot wizard to create a snapshot of a virtual disk at a specified time and frequency. If specified, the wizard deletes any previous snapshots with the same name.

Delete Snapshot

Use the Delete Snapshot wizard to delete a snapshot at a specified time. If the specified snapshot does not exist, nothing is deleted.

Restore From Snapshot

Use the Restore From Snapshot wizard to restore a virtual disk from a snapshot. If specified, the wizard deletes the the old virtual disk and its snapshot after restoring the virtual disk. However, if the virtual disk had more than one snapshot, the old virtual disk cannot be deleted. This is because the restore operation delete the restored snapshot first and the virtual disk second. The virtual disk delete operation fails if the virtual disk still contains other snapshots. The activities are tracked in the event log.

Snapshot For Backup

Use the Snapshot For Backup wizard to create and delete a snapshot for a backup.

Snapshot Watchdog

Use the Snapshot Watchdog wizard to schedule a task that monitors the resource utilization of a snapshot at a specified frequency. You can specify that the snapshot be deleted if the resource utilization exceeds preset values.

Workday Snapshot

Use the Workday Snapshot wizard to create a snapshot of a virtual disk at a specified time (for example, at the start of the work day) and to delete the snapshot after a specified time (for example, at the end of the work day).

Managing SVR Policies

10

This appendix describes how to manage SVR policies:

Policy overview

Policies are system settings that help you manage SVR resources. The SVR policies are stored in the Windows Registry and are preconfigured to maximize performance. You can modify these default policies for pools, virtual disks, and snapshots for your environment.

You can configure the policies to apply to the entire SVR environment or set them for individual pools. For example, you set a policy that reserves a percentage of a pool's free space exclusively for snapshots. You can apply this policy to all pools on your system or to selected pools.

If you use Policy Editor on a cluster system to edit policy keys, the policy keys on both nodes are updated simultaneously.

Using Policy Editor

This section describes how to use Policy Editor:

- Overview
- Access
- Options

Overview

Policy Editor enables you to manage SVR policies efficiently. You can set policies to govern the creation and use of SVR resources. Policy Editor provides instructions to make registry changes easy and prevent accidental changes.

Any changes you make to the SVR registry settings using Policy Editor are effective immediately. The exception is the `SSDeleteSnapshot` policy, which requires that you restart the Lifeguard service.

Caution: HP strongly recommends that you use Policy Editor (not the Windows Registry) to edit policy settings. If an error occurs in the Windows Registry, your computer may not function properly. You must be prepared to restore the registry if an error occurs.

Access

You can access Policy Editor using either of the following methods:

- Using Replication Manager
- Using the Properties window

Using Replication Manager

To access Policy Editor from Replication Manager:

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

   The Replication Manager window opens.

2. Do one of the following:
   - Right–click in the scope pane and select **All tasks > Policy Editor**.
   - Click the **Policy Editor** icon on the Replication Manager toolbar

   The **Policy Editor** window opens (Figure 45).

The **Policy Manager** window displays the following policies:

- **Default**—Default policies that are preconfigured to follow best practices for managing SVR.
- **Global**—System-wide policies that are applied to all new pools.
- **Pool**—Policies that apply to specific pools only.

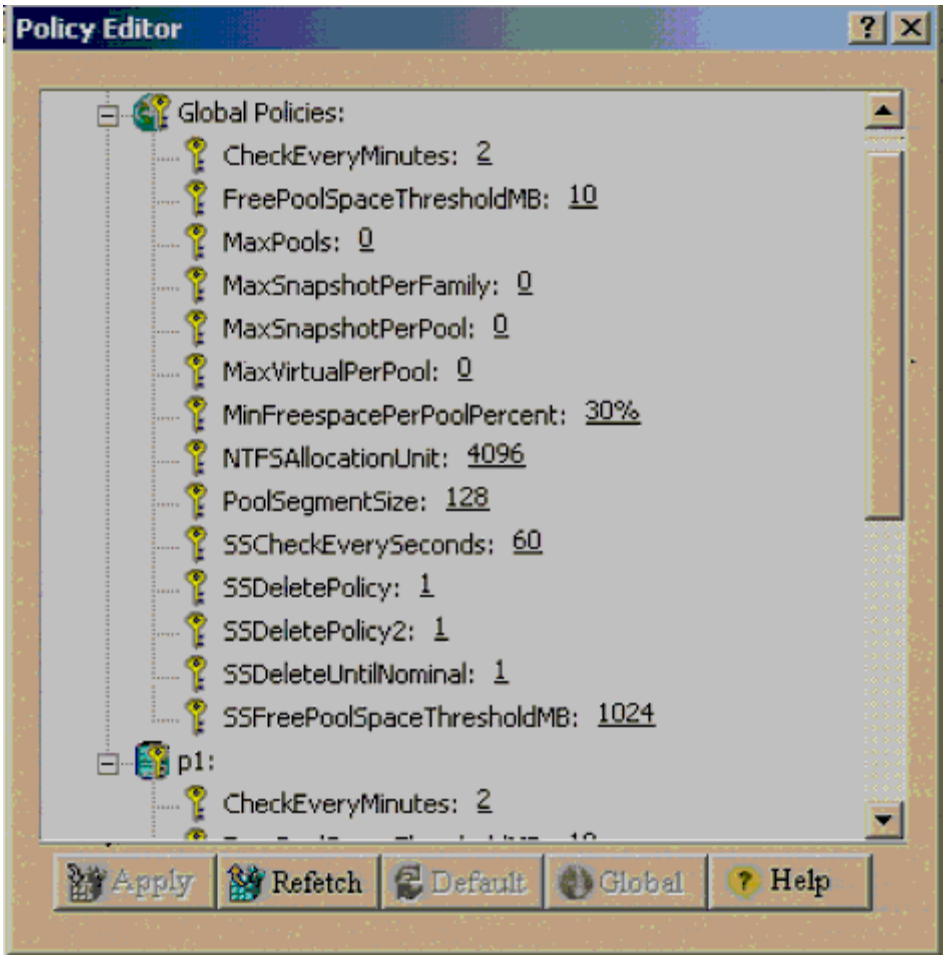

Figure 45: Policy Editor window

Using the Properties window

To access Policy Editor from the pool Properties window:

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

   The Replication Manager window opens.

2. Right–click the pool and do either of the following:

   - Select **All tasks > Policy Editor**.
   - Select **Properties** and then click the **Policy Editor** tab.

   The Policy Editor window opens (Figure 45).

The Policy Editor window displays the following policies:

- **Global**—System-wide policies that are applied to all new pools.
- **Pool**—Policies that apply to specific pools only.

Options

The following options are available in the Policy Editor window:

- **Reset to Default**—Resets a policy to the default setting. You can apply the default setting to one pool policy or all pool policies.
- **Reset to Global**—Customizes the default policies for your environment. After you modify a global policy, all new pools inherit the Global policy values. You can select **Reset to Global** for one pool policy or all pool policies.
- **Apply**—Applies the policy changes you make. You are asked to confirm this action before the new policy settings are applied. After the settings are changed, you cannot revert back to the previous settings.
- **Refetch**—Retrieves the current policy settings.
- **Close**—Close the Policy Editor window.
- **Policy Alert**—View suggestions for settings policy values. These alerts ensure that you set the policy values correctly.
- **Context-sensitive help**—Provides brief descriptions of SVR policies. Click the question mark in the top-right corner of the Policy Editor window and then select the appropriate policy.

Setting Policies

To set a policy:

1. Open Policy Editor.
2. Do one of the following:
   - If the policy you are modifying affects all pools, locate the appropriate policy under **Global Policies**.
   - If the policy modification affects one pool, locate the appropriate policy under the specific pool.
3. Select the policy name to activate the value box (Figure 46).
4. Enter the value for the policy.
5. Click **Apply**.

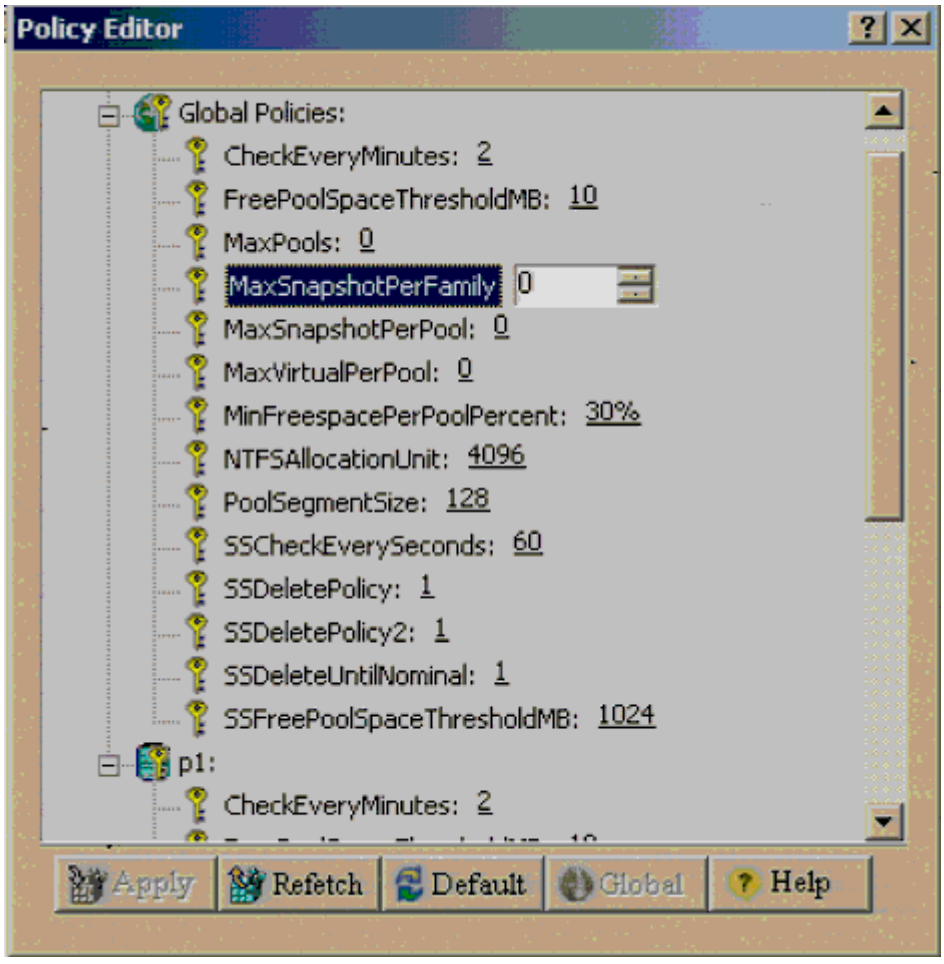

Figure 46: Selecting a policy

Setting pool policies

This sections describes the policies you can set for pools. It describes the following topics:

- MaxPools
- MinFreespacePerPoolPercent
- PoolSegmentSize

MaxPools

The `MaxPools` policy limits the number of pools that can be created on the system. There is no limit to the number of pools on a system unless you set this policy. This policy is not enforced in a cluster; the number of pools is unlimited.

When the number of pools equals the policy value, you cannot create additional pools.

Note: The number of pools includes failed pools.

MinFreespacePerPoolPercent

The `MinFreespacePerPoolPercent` policy reserves a percentage of the pool's free space exclusively for snapshots. You cannot create or grow a virtual disk if it would cause the pool's free space to fall below the specified threshold. The default setting of 30% ensures you do not use all available free space when growing a virtual disk, which may result in lost snapshots and data.

Although you can set the policy lower, HP strongly recommends that you do not decrease this minimum amount.

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

Examples

You have a pool that is 200 GB and the `MinFreespacePerPoolPercent` policy is set to 30%, which equals 60 GB. You create a virtual disk of 90 GB, leaving 110 GB of free space. You cannot create a second virtual disk of 90 GB, because it would leave 20 GB of free space. If you do not need snapshots, you can create a second virtual disk of 50 GB and still maintain the 30% threshold.

In another example, you have a pool that is 150 GB and the `MinFreespacePerPoolPercent` policy is set to 30%, which equals 45 GB. The pool contains a virtual disk of 70 GB and one snapshot using 10 GB, leaving 70 GB of free space. You cannot grow the virtual disk by 70 GB, because it would use all the space in the pool. You can grow the virtual disk by 25 GB, which would leave 45 GB of free space.

PoolSegmentSize

The `PoolSegmentSize` sets the size of the pool upon creation. The default segment size is 128 KB, which enables you to 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 becomes the default segment size during pool creation. However, SVR overrides this policy default if you create a pool and add a large storage unit that requires a larger segment size.

The `PoolSegmentSize` policy applies to all pools. You cannot set this policy for individual pools.

Setting virtual disk policies

This sections describes the policies you can set for virtual disks. It describes the following topics:

- MaxVirtualPerPool
- NTFSAllocationUnit

MaxVirtualPerPool

`MaxVirtualPerPool` sets the maximum number of virtual disks per pool. The policy limits virtual disk creation. SVR allows a maximum of 8 virtual disks per pool; therefore, only values of 8 or less are valid.

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

NTFSAllocationUnit

`NTFSAllocationUnit` changes the default NTFS allocation unit size for a virtual disk. By default, when formatting virtual disks, SVR 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 modify this policy, the revised value becomes the new default. This new default value is used when:

- You format a virtual disk using Replication Manager and set the allocation unit size to **Default**.
- You format a disk using `SnapMgr` and do not specify an allocation unit size.

To override the `NTFSAllocationUnit` policy, you can:

- Select a size another than Default in the Replication Manager virtual disk wizard.
- Select a size another than Default using `SnapMgr`.
- Use Windows Explorer to format the disk.

Note: This policy has no effect on other Windows disks—only on virtual disks formatted using SVR.

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.

Setting snapshot policies

This sections describes the policies you can set for snapshots:

- MaxSnapshotPerPool
- MaxSnapshotPerFamily

MaxSnapshotPerPool

`MaxSnapshotPerPool` sets the maximum number of snapshots that you can create in a pool. SVR allows a maximum of 12 snapshots for each virtual disk and 8 virtual disks for each pool. Therefore, only values of 96 or less are valid.

SVR checks the `MaxSnapshotPerPool` and `MaxSnapshotPerFamily` policies simultaneously. The policy limit that is reached first is executed.

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

MaxSnapshotPerFamily

`MaxSnapshotPerFamily` limits the maximum number of snapshots that you can create in a family. A virtual disk and its child snapshots are known as a family. SVR allows a maximum of 12 snapshots for each virtual disk (family). Therefore, only values of 12 or less are valid.

SVR checks the `MaxSnapshotPerPool` and `MaxSnapshotPerFamily` policies simultaneously. The policy limit that is reached first is executed.

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

Setting SVR Lifeguard policies

This section describes the SVR Lifeguard service and the policies you can set for it:

Using SVR Lifeguard

SVR Lifeguard is a system service that monitors a pool's free space and deletes snapshots when the pool becomes full. This service is installed and started automatically. Lifeguard first checks each drive letter on a system to determine if it is a virtual disk or snapshot. If it is either, Lifeguard checks the disk's pool free space every 60 seconds.

SVR 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 for each snapshot that is deleted. See "SVR Lifeguard events" on page 150 for more information.

Using Lifeguard on Microsoft Exchange pools

If you install Microsoft Exchange on a virtual disk, 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 determine if it is a virtual disk. If so, Lifeguard checks the disk's pool free space every two minutes.

SSDeletePolicy

`SSDeletePolicy` sets the free space limit for a pool. When the free space in a pool drops below the free space limit, Lifeguard deletes snapshots, beginning with the oldest, until free space reaches that limit. The default value is 1024 MB.

If `SSDeletePolicy` is not executed, `SSDeletePolicy2` specifies that Lifeguard should delete snapshots with the largest Delspace to free up space in the pool.

For Microsoft Exchange pools, Lifeguard shuts down Microsoft Exchange 5.5 when the storage pool has 10 MB left. You can review the Application Event log to determine which pools are full. After you add more space to the appropriate pools, restart Microsoft Exchange.

Note: Automatic shutdown is supported on Microsoft Exchange 5.5 only. If you are using Microsoft Exchange 2000, HP recommends that you set the deletion policies to prevent pools from running out of disk space.

SSCheckEverySeconds

`SSCheckEverySeconds` sets the frequency with which Lifeguard checks pools for free space. The default frequency is 60 seconds. This policy applies to all pools and cannot be set separately for individual pools.

Note: Setting the value to 0 does not disable Lifeguard monitoring.

CheckEveryMinutes

`CheckEveryMinutes` sets the frequency with which Lifeguard checks Microsoft Exchange storage pools for free space. The default interval is 2 minutes. This policy applies to Microsoft Exchange pools only.

Note: Setting the value to 0 does not disable Lifeguard Exchange monitoring.

SSFreePoolSpaceThresholdMB

`SSFreePoolSpaceThresholdMB` sets the pool free space threshold. When pool free space falls below this threshold, Lifeguard deletes snapshots. The default threshold is 1024 MB. This policy is dependent on the `SSDeletePolicy` or `SSDeletePolicy2` policy.

FreePoolSpaceThresholdMB

`FreePoolSpaceThresholdMB` sets the threshold of free space below which Lifeguard shuts down Microsoft Exchange 5.5. The default threshold is 10 MB. This policy applies to Microsoft Exchange pools only.

Note: Setting this value to 0 will disable Exchange monitoring.

SSDeleteUntilNominal

`SSDeleteUntilNominal` specifies the number of snapshots to be deleted when the pool's free space reaches the threshold specified by the `SSFreePoolSpaceThresholdMB` policy. If you set the value to 0, one snapshot is deleted at each interval. If you set the value to 1 or higher, that number of snapshots is deleted at each interval. This policy is dependent on the `SSDeletePolicy` or `SSDeletePolicy2` policy.

Troubleshooting SVR

11

This chapter describes how to troubleshoot issues that may occur when using SVR. It describes the following topics:

General troubleshooting questions

Be prepared to answer the following questions that apply to your situation when contacting HP Support for assistance. You can also use these questions to determine why a pool did not come online in a clustered environment.

- What version of SVR are you using?
- Do you have other resource dependencies defined within the pool group?
- What caused the node to fail? Was it related to SVR?
- What type of storage controllers (for example, G80, EVA, or MSA) do you have?
- Did you make any controller firmware upgrades to the hardware?
- Were there any errors logged at the controller level?
- If you shut down the node when it was in the 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 failed and the pool moved to the other node, were you logged in to the other node? (This applies to SVR version 3.x.)

Incorrect drive letters after upgrade or cluster failover

Issue

Virtual disks and snapshots may lose their drive letter mappings when:

- You upgrade SVR.
- A pool fails over in a cluster.

If the drive letter that is mapped to a virtual disk or snapshot is not available on the second node, SVR unmounts the drive letter on the second node and reassigns it to the volume that failed over. However, if the volume that failed over has the same drive letter as the system disk of the second node, then the drive letter is not mapped.

Resolution

You can use the `Restore Drives` utility to restore drive letters to virtual disks and snapshots. The utility is installed with SVR and is located in the following directory:

```
Program Files\Hewlett-Packard\OpenView Storage Virtual
Replicator 4.1
```

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

```
C:\Program Files\Hewlett-Packard\OpenView Storage Virtual
Replicator 4.1\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, the Restore Complete message is displayed. You can verify that the drive letters were mapped correctly in the Replication Manager window.

Lost delayed write errors for snapshots

Issue

If you receive error messages that delayed writes have been lost while you are using snapshots, it may indicate that the pool is full.

Resolution

To determine if the pool is full:

1. Right–click **Replication Manager** in the scope pane and select **Connect to another computer**.

   The Select Computer window opens.

2. See "Creating a temporary remote connection" on page 32 for instructions about connecting to a remote computer.

3. View the Free column in the results pane to determine how much free space is available for all pools on the selected computer.

   Do any pools show 0 MB in the Free column?

   - If yes, go to step 4.
   - If no, the problem you are experiencing is 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. Do one of the following:

   - Add a storage unit to the pool:
     - Right–click the pool and select **Properties**.
     - Select the Storage Units tab and click **Add**.
     - Select a storage unit and click **Add**.
     - Click **OK** twice.

     See "Using storage units in pools" on page 47 for more information.
   - Delete a snapshot or virtual disk from the pool:
     - Right–click the virtual disk or snapshot and select **Delete**.
     - Click **OK** to confirm that you want to delete the disk.

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

Applications return failed write errors

Issue

You are using snapshots and applications return errors when you try to write data to a disk.

Resolution

To determine if the pool is full:

1. Right–click **Replication Manager** in the scope pane and select **Connect to another computer**.

   The Select Computer window opens.

2. See “Creating a temporary remote connection” on page 32 for instructions about connecting to a remote computer.

3. View the Free column in the results pane to determine how much free space is available for all pools on the selected computer.

   Do any pools show 0 MB in the Free column?

   - If yes, go to step 4.
   - If no, the problem you are experiencing is 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. Do one of the following:

   - Add a storage unit to the pool:
     - Right–click the pool and select **Properties**.
     - Select the Storage Units tab and click **Add**.
     - Select a storage unit and click **Add**.
     - Click **OK** twice.

     See “Using storage units in pools” on page 47 for more information.
   - Delete a snapshot or virtual disk from the pool:
     - Right–click the virtual disk or snapshot and select **Delete**.
     - Click **OK** to confirm that you want to delete the disk.

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

Deleting a virtual disk accidentally

Issue

You delete a virtual disk accidentally. When you delete a virtual disk, SVR does not overwrite the physical blocks that store the disk's data.

Note: When you delete a virtual disk, the first few blocks of data are erased. You must rebuild these blocks. Contact support for assistance.

Resolution

Follow these steps *immediately* to save the data:

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

   Using the New Virtual Disk wizard, map a drive letter to the new virtual disk but do not format it. Click **No** when you are prompted to format the disk.

2. In Windows Explorer, double–click the icon for the new virtual disk. You may see the files and folders of the virtual disk you deleted accidentally. The new virtual disk is using the same disk blocks as the one you deleted accidentally.

Note: If a message is displayed indicating that the drive is inaccessible and does not contain a file system, you must restore the data from your backup.

Pool free space decreases by more than the amount of written data

Issue

When a disk contains a snapshot and you write data to either the snapshot or its parent disk, the pool space may decrease by more than the amount of written data.

Explanation

The decrease occurs because the data is written in defined chunks, called *segments*. You define the segment size when you create a pool. (See “Creating pools” on page 42.) If a segment contains data modified by the write, the entire segment is copied to the disk. For example, if you write 2 KB of data and 1 KB is in one segment and 1 KB is in another segment, the pool free space decreases by 64 KB (two 32 KB segments). Generally, the disk space used for a write is approximately the same size as the write. You should see a difference only in unusual situations, for example, when running a test program that writes small amounts of data randomly over the surface of the disk.

Cannot reformat a virtual disk

Issue

When you reformat a virtual disk using Replication Manager, either the `Format Complete` or `Succeeded` message is displayed; however, the disk has not been reformatted.

Resolution

You cannot reformat a virtual disk if another application is accessing the virtual disk. For example, you cannot reformat a virtual disk if your or another user is accessing that disk in Windows Explorer or a command prompt window.

Ensure that users are not accessing the virtual disk before you use Replication Manager to reformat it.

Reconstructing a pool

Issue

A pool has failed.

Resolution

To reconstruct a pool:

1. Determine 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 to restore saved data to the new virtual disks.

Pools do not come online

Issue

If the pools do not come online, view each pool to determine which pool is causing the problem.

Resolution

If the problem is due to a deletion in the registry, you must uninstall and reinstall SVR to update the registry:

1. Power down one node.
2. From the desktop of the other node, select **Start** > **Settings** > **Control Panel** > **Add/Remove programs**.
3. Select **Virtual Replicator** and click **Remove**.
4. Power up the node and open **Disk Management**. Verify that the disks in the pool are available.

Note: If you do not see all of the disks that belong in the pool, SVR cannot rebind the pool. The problem is not related to SVR and is most likely a storage problem.

---

5. Reinstall SVR.
6. To resolve namespace errors, enter the following command:

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

   This command updates the registry using the online pools.
7. If applicable, issue the `RestoreDrives` command.

   The pools that come back online enable the metadata on the drive to rebuild the pool.
8. Repeat this procedure on the other node.

Note: This procedure requires approximately four reboots and one hour to complete.

Experiencing prolonged startup times

Issue

If you experience prolonged startups or reboots, it may be that SVR pools require more time to *rebind*. Rebinding occurs during system startup and enables SVR pools to reestablish the address links between parent disks and snapshots and load those links into memory.

Typically, rebinding requires little time and does not affect performance. However, if the differences between the snapshots and the parent disks are significant, the time needed to rebind increases. Differences occur when you make changes to the parent disk and segments are copied out to the snapshots.

Copy-outs are required when:

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

Resolution

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

- Avoid saving 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-outs.

SVR names are already being used

Issue

You create a pool, virtual disk, or snapshot and a message is displayed, indicating that the name is already being used.

This can occur if there are namespace errors. When you create a pool or add virtual disks, snapshots, and logical drives to a pool, SVR tracks their names in the Windows registry. If a pool fails, Replication Manager does not display that pool or the items contained in it but the names are still listed in the Windows registry.

Resolution

To resolve namespace errors, enter the following command:

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

This command updates the registry using the online pools.

Tracking errors

Use the Windows Event Viewer to track errors.

There are two types of events that produce error, warning, and informational messages:

- SVR driver events (`hpqvrbus`)
- SVR Lifeguard events (`swvrmon`)

The tables in the following sections provide details about the events that SVR writes to the Event Viewer logs.

SVR driver events

Table 12 describes the `hpqvrbus` messages written to the Event Viewer system log.

Table 12: SVR driver event messages

Severity	MessageId	Event source	Event description
Informational	1	`hpqvrbus`	The SVR driver has started.
Error	2	`hpqvrbus`	The SVR driver failed at initialization.
Error	3	`hpqvrbus`	There is insufficient memory.
Error	4	`hpqvrbus`	SVR cannot create device %2, where %2 is the virtual disk or snapshot name that could not be created or initialized.
Error	5	`hpqvrbus`	SVR could not read or write to a disk in the pool.
Error	6	`hpqvrbus`	SVR could not create or initialize disk %2, where %2 is the virtual disk or snapshot name that could not be created or initialized.
Error	7	`hpqvrbus`	SVR could not get pointer to device %2, where %2 is the disk device name that could not be found.
Warning	8	`hpqvrbus`	Free disk space in the pool %2 (where %2 is the pool name) is less than 30%. Delete any unneeded snapshots.
Error	9	`hpqvrbus`	Free disk space in the pool %2 (where %2 is the pool name) is less than 5%. Delete any unneeded snapshots.

SVR Lifeguard events

The SVR Lifeguard service monitors pools and generates `swvrmon` events. These events are recorded in the Event Viewer application log.

Table 13: Lifeguard event messages

Severity	MessageId	Event source	Event description
Informational	1	`swvrmon`	The SVR Lifeguard service has started.
Informational	2	`swvrmon`	The SVR Lifeguard service has stopped.
Error	3	`swvrmon`	The service handler has not been installed.
Warning	4	`swvrmon`	SVR pool %2 (where %2 is the pool name) has reached Microsoft Exchange Server shutdown threshold for partition %3 (where %3 is the drive letter of the virtual disk or snapshot): of < %4 MB free in the pool (where %4 is the minimum free pool space threshold).
Error	5	`swvrmon`	The SVR Lifeguard service will now stop Microsoft Exchange Server.
Error	6	`swvrmon`	Error %2 (where %2 is the error number) stopping service %3 (where %3 is the Microsoft Exchange service name).
Informational	7	`swvrmon`	The SVR Lifeguard service has completed shutdown of Microsoft Exchange Server.

Table 13: Lifeguard event messages

Severity	MessageId	Event source	Event description
Warning	8	`swvrmon`	SVR pool %2 (where %2 is the pool name) has reached the snapshot deletion policy threshold for partition %3 (where %3 is the drive letter of the virtual disk or snapshot): of < %4 MB free in the pool (where %4 is the minimum free pool space threshold—typically 10 MB of disk space).
Error	9	`swvrmon`	The SVR Lifeguard service will now execute snapshot delete policy %2 (%3) (where %2 is the policy type and %3 is the policy number) for pool %4 (where %4 is the pool name) on snapshot %5 (where %5 is the snapshot name). The policy number (%3) and type (%2) can be any of the following: 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 13: Lifeguard event messages

Severity	MessageId	Event source	Event description
Error	10	`swvrmon`	The SVR Lifeguard service will now execute second-chance snapshot delete policy %2 (%3) (where %2 is the policy type and %3 is the policy number) for pool %4 (where %4 is the pool name) on snapshot %5 (where %5 is the snapshot name). The policy number (%3) and type (%2) can be any of the following: 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

index

R

S

T

U

V

W