

DATA PROTECTION USING PREMIUM FEATURES

POWERSHIELD™ MD3000 AND MD3000i

Dell | Product Group



CONTENTS

DELL™ POWERVAULT™ MD3000 AND MD3000i PREMIUM FEATURES	3
ACTIVATING PREMIUM FEATURES	3
VIRTUAL DISK SNAPSHOT	
Using Snapshots	4
Creating a Snapshot	5
Snapshot Repository Virtual Disk	5
Re-creating a Virtual Disk Snapshot	6
Virtual Disk Snapshot and the Microsoft® Volume Shadow Copy Service (VSS)	6
VSS Architecture	7
Using VSS to Create Snapshots	8
VIRTUAL DISK COPY	8
Using Virtual Disk Copy	9
Creating a Virtual Disk Copy	9
Re-copying a Virtual Disk Copy	10
COMBINING THE VIRTUAL DISK COPY AND VIRTUAL DISK SNAPSHOT PREMIUM FEATURES	10
Creating a Virtual Disk Copy from a Snapshot	10
Restoring Snapshot Data to the Source Virtual Disk using Virtual Disk Copy	11
CONCLUSION	12
Benefits of Virtual Disk Snapshot	12
Benefits of Virtual Disk Copy	12
Benefits of Combining Virtual Disk Snapshot and Virtual Disk Copy	12
CONTACT INFORMATION	13
APPENDIX A: COMPARING A SNAPSHOT TO A VIRTUAL DISK COPIES	13
REFERENCES	14
TABLES	
TABLE 1: COMPARISON OF VIRTUAL DISK SNAPSHOT AND VIRTUAL DISK COPY	13
FIGURES	
FIGURE 1: THE VIEW/ENABLE PREMIUM FEATURES DIALOG	3
FIGURE 2: COPY-ON-WRITE	4
FIGURE 3: VSS COMPONENTS	7
FIGURE 4: CREATING A VIRTUAL DISK COPY	8
FIGURE 5: A SNAPSHOT VIRTUAL DISK AS THE SOURCE FOR CREATING A VIRTUAL DISK COPY	10
FIGURE 6: SELECTING THE SOURCE VIRTUAL DISK	11
FIGURE 7: VIRTUAL DISK COPY PROCESS FOR SNAPSHOT VIRTUAL DISK TO SOURCE VIRTUAL DISK	12

SECTION 3

VIRTUAL DISK SNAPSHOT

A *snapshot virtual disk*, also called simply a snapshot, is a point-in-time image of a virtual disk. A snapshot is functionally equivalent to a complete physical copy. However, a snapshot is created almost instantaneously and requires less disk space than the source virtual disk. To a host server a snapshot appears and functions just as a standard virtual disk does. A snapshot can be assigned to a host server to be read, written to, or copied from.

A snapshot uses copy-on-write technology to maintain the snapshot while minimizing disk utilization. When a snapshot of a virtual disk is taken, the storage array controller suspends I/O to the source virtual disk momentarily while it creates a repository virtual disk. The repository virtual disk stores snapshot copy-on-write data. When data on the source virtual disk is modified, a copy-on-write occurs, copying the data that is to be modified into the repository virtual disk. This repository virtual disk combined with the source virtual disk creates the logical snapshot virtual disk. Because the only data that is stored in the repository virtual disk is data that has changed after the creation of the snapshot, a snapshot uses less disk space than a full physical copy. The default size for the repository virtual disk is 20 percent of the source virtual disk capacity.

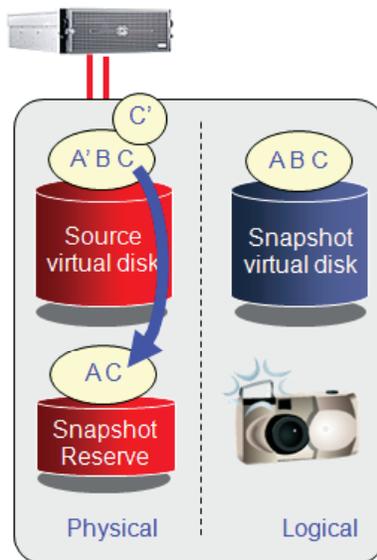


Figure 2: Copy-on-Write

Each virtual disk can support up to four snapshots, with a maximum of 128 snapshots per storage array.

USING SNAPSHOTS

Snapshots provide many benefits for data backup and data manipulation.

Continued data processing during backups

The Virtual Disk Snapshot premium feature enables offline backups of production data without disrupting access to that data. Snapshot creation requires that the application or file system be quiesced only briefly. After the snapshot is created, access to the original data is again available for business processing, and the point-in-time image of the data is ready for backup.

The Dell Modular Disk Storage Manager software allows snapshots to be mapped to a separate backup server. This capability is designed to reduce the performance impact to the application server by placing the burden of the backup processing on an alternate server.

Data protection through easy, fast restoration of data

The time and cost to back up data is a major consideration, but more important is the ability to restore data and to restore it quickly. Tape restorations can result in a day or more of lost time. In contrast, the ability to create multiple snapshots in a single day allows access to the most recent data possible for recovery, helping to minimize recovery time.

Application Testing

Application testing is a necessary practice in 24x7 IT environments. New applications and application patches generally must be tested before they can be put into production. Snapshots are designed to allow for the rapid creation of a point-in-time image to perform these tests against. Testing on these non-production images essentially eliminates the risk of lost production data or lost data access. It can also save system administrators time by eliminating the need to restore production data that might become corrupted or lost due to an incompatible application or patch.

Data Analysis

Snapshots can provide point-in-time images for data analysis. The images can be created almost instantaneously, allowing for the data analysis to occur on the snapshot rather than the production virtual disk. Any performance impact of the data analysis will be incurred on the snapshot rather than on the production virtual disk. Additionally, the smaller disk space requirement of a snapshot—which is by default 20 percent of the source virtual disk—saves disk space for the creation of additional test or analysis environments.

CREATING A SNAPSHOT

A snapshot can be created using the Dell Modular Disk Storage Manager software's graphical user interface or through the command line interface (CLI). The CLI has the added benefit of command scripting, which allows for task automation and scheduling.

NOTE: Before creating or re-creating a snapshot, it is recommended to stop all I/O activity to help ensure that you capture an accurate point-in-time image of the source virtual disk. You also should run the SMrepassist utility on the source virtual disk. SMrepassist (replication assistance) is a host-based utility for Windows® that is designed to ensure that all the memory-resident data for the file system is flushed to disk. In addition, SMrepassist updates disk signatures to help prevent conflicts when the snapshot is assigned to the same host as the source virtual disk.

Refer to the appropriate Dell™ PowerVault™ documentation for procedures on your chosen method for creating a snapshot virtual disk.

Snapshot Repository Virtual Disk

For every snapshot that is taken, a repository virtual disk is created to hold the copy-on-write data for the snapshot. There is a one-to-one relationship between a snapshot and its associated repository virtual disk. The repository virtual disk is not accessible for host I/O. The repository virtual disk is only used internally to manage the snapshot point-in-time image.

The required size of the repository depends on the size of the source virtual disk and the write patterns that occur on the source after the snapshot is established. If the snapshot is created as a point-in-time image for data backup, a default repository size of 20 percent or less of the source virtual disk is usually sufficient. If the snapshot is created for functions other than backup, such as application testing or data manipulation, a larger repository might be needed. The repository size depends on the rate at which data changes in the source virtual disk, the snapshot virtual disk, and on the length of time that the snapshot remains active.

A best practice is to create the repository with a default value of 20 percent, and then use the following procedure to estimate the life expectancy of the snapshot and its associated repository.

1. Record the creation date and the creation time for the snapshot.
2. After a period of operation in your environment, record the repository virtual disk capacity used, the available repository capacity remaining, and the current time.
3. Subtract the snapshot creation time from the current time to determine the elapsed time (length of time the snapshot has been active).
4. Use the following equation to calculate the amount of time that the snapshot repository will be available for copy-on-write data before reaching its capacity:

$$\text{Total time available} = (\text{elapsed time} \times \text{available capacity}) / \text{used capacity}$$

Based on the results, you can decide how much to increase the capacity of the repository virtual disk. If the repository will become 100 percent full during the expected life time of the snapshot, you should increase the capacity of the repository virtual disk. A snapshot's repository virtual disk can be dynamically expanded at any time.

To alert yourself of an impending "repository full" condition, configure a warning threshold for repository free capacity when you are creating the snapshot. The default and the recommended value for the threshold is 50 percent. This value should provide enough advanced warning to allow you time to assess the repository virtual disk use and to expand the capacity if required. By default, the snapshot will be failed if the repository becomes completely full. The option to fail writes to the source virtual disk is also available. This option allows the snapshot to remain active, and writes to the source virtual disk resume after the repository capacity has been increased.

You can locate the snapshot repository virtual disk in any disk group. However, to simplify physical disk management, locate the repository virtual disk in the same disk group as the source virtual disk whenever possible. This ensures that if you must move the physical disks associated with the disk group to another storage array, all the virtual disks associated with the snapshot reside on that group and the snapshot remains accessible.

Re-creating a Virtual Disk Snapshot

When you are finished working with a snapshot, you have the option of disabling or deleting it. Deleting a snapshot deletes the repository virtual disk and associated data. Disabling a snapshot saves the snapshot and its associated repository for later use. Using a disabled snapshot to re-create a snapshot takes less time than creating a new snapshot. For example, if you need to create daily backups of the same source virtual disk, rather than creating a new snapshot every day, you simply re-create the existing disabled snapshot. The Re-create Snapshot option is extremely beneficial in environments where frequent backups or frequent manipulations of the same data set occur.

NOTE: All data previously written to the snapshot is overwritten when a snapshot is re-created.

VIRTUAL DISK SNAPSHOT AND THE MICROSOFT® VOLUME SHADOW COPY SERVICE (VSS)

The Microsoft Volume Shadow Copy Service, also called Volume Snapshot Service (VSS), is a storage management interface for Microsoft® Windows Server® 2003 and Microsoft Windows Server 2008 operating systems that allows you to create snapshot data images. VSS enables your storage array to interact with third-party applications that use the VSS Application Programming Interface (API) to create snapshots that are consistent by ensuring that all memory resident application data is flushed to the

storage array before the snapshot is taken. Typical uses for VSS include data backup and restoration and data transport (only supported through a VSS hardware provider) for testing, data mining, or backup. Dell supplies a VSS hardware provider on the PowerVault MD Resource CD to enable VSS functionality on the PowerVault storage arrays.

VSS Architecture

The VSS architecture consists of three components that interact to help create a consistent snapshot. Overall coordination of the snapshot creation process is under VSS framework in the operating system.

REQUESTOR – A *requestor* is a third-party application, usually a backup application that requests VSS to perform a shadow copy (snapshot) of a particular virtual disk. The requestor initiates the VSS process to create a snapshot.

WRITER – A *writer* is an application provided module which takes application-specific action to ensure that its data store on the virtual disk is consistent (memory resident data is flushed) before the snapshot is taken. Applications such as Microsoft Exchange Server 2003 (and later) and Microsoft® SQL Server® 2005 (and later) are examples applications that incorporate VSS writers.

PROVIDER – A *provider* is the interface to the storage array used by the VSS framework to request the storage array to create a snapshot. The provider can be a hardware device specific, such as the Dell Modular Disk VSS provider, or a software provider like the one included in Windows Server 2003 or 2008. The VSS framework attaches to the provider and uses it to coordinate the creation of snapshots on the storage array.

Figure 3 shows the interacting components in a typical VSS configuration.

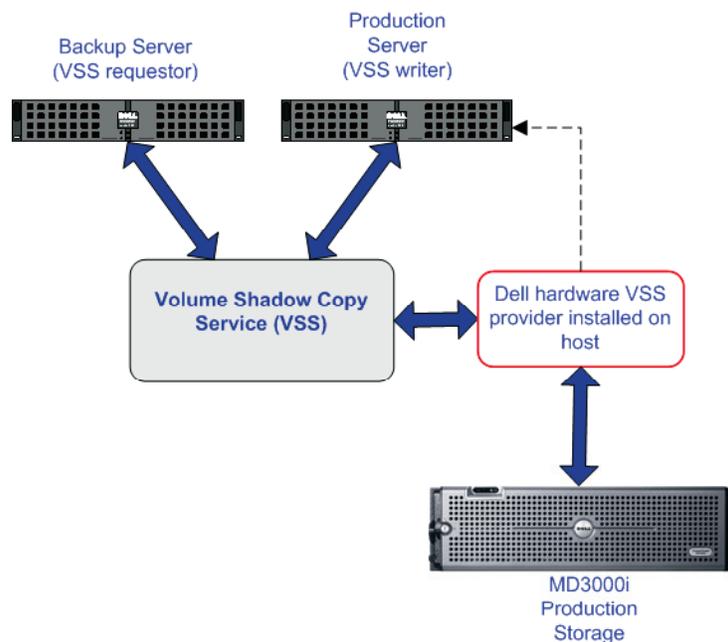


Figure 3: VSS Components

Using VSS to Create Snapshots

The most common use of VSS is in backup applications. Most third-party backup applications today provide VSS requestor functionality. Used in conjunction with a VSS aware application (Microsoft SQL Server, Microsoft Exchange) and the Dell hardware VSS provider, backup applications can automatically create snapshots on the storage array and back them up according to the user-defined backup policy. Because VSS manages the entire process, there is no need to manually quiesce or prepare the data for snapshot creation. Applications that are integrated with VSS automatically perform these operations. Pre-defined policies are designed to help reduce the time and the complexity associated with data backups and thus help ensure consistent backups that can be restored quickly.

The Dell™ PowerVault™ MD3000 and MD3000i storage arrays and the Dell hardware VSS provider also make it possible to create a snapshot of the source data on one server, and then to import that snapshot onto another server. This process, known as snapshot transport, is accomplished in minutes, regardless of the size of the data. This transport capability reduces the performance impact to the application server by placing the burden of the backup processing on an alternate server.

SECTION 4 VIRTUAL DISK COPY

Virtual Disk Copy creates a physical copy, or a clone, of a source virtual disk. A virtual disk copy appears and functions as a standard virtual disk. A virtual disk copy can be addressed by the host and can be read, written to, or copied from. Virtual Disk Copy is a background operation with user-defined priority settings, enabling administrators to minimize either copy time or the overall I/O impact to the storage array. As Virtual Disk Copy is based on the storage array controller, creating a virtual disk copy requires no host interaction or CPU cycles — minimizing the impact to applications and the storage infrastructure.

The virtual disk copy relationship is established between a source and a target virtual disk. To help ensure data integrity, all I/O to the target virtual disk is suspended during the procedure. This is because the state of data on the target virtual disk is inconsistent until the procedure is complete. Also, all updates to the source virtual disk are suspended to help prevent chronological inconsistencies from being created on the target virtual disk. After the Virtual Disk Copy operation is complete, the target virtual disk automatically becomes read-only to the hosts unless the default write protection setting is altered. In order to preserve the data on the target virtual disk, it is best to keep the write protection setting set to Read-Only.

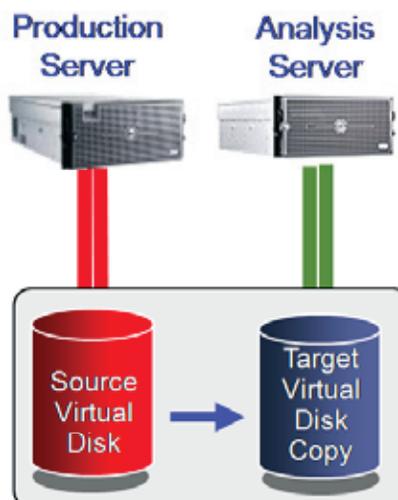


Figure 4: Creating a Virtual Disk Copy

When Virtual Disk Copy is enabled, the PowerVault's MD3000 storage array and MD3000i storage array each support up to 255 virtual disk copies and up to eight concurrent Virtual Disk Copy operations.

USING VIRTUAL DISK COPY

Virtual Disk Copy offers many benefits for data backup and manipulation.

Data Backup and Restoration

The virtual disk copy procedure can be leveraged into a high-throughput backup solution providing disk-based backup. The storage array manages the data transfer internally, intelligently balancing backup-related traffic against production traffic on the drive channels. This is designed to minimize the impact on production processing and to isolate all traffic within the array itself, with the goal of eliminating the requirement for additional external I/O channels for backup purposes. Because the operation is carried out entirely within the array, the backup process is designed to be completed as efficiently as possible and to outperform conventional tape-based systems, server-to-tape systems, and even other external disk-based systems.

The data can be recovered almost instantly by simply unmapping the source production virtual disk, and mapping the backup virtual disk in its place. In this case, the roles of the virtual disks change: The previous backup virtual disk becomes the new production virtual disk, and the old production virtual disk can be re-used as a target for future backups.

Virtual Disk Relocation

Virtual disk requirements can and do change over time. The overall required capacity might increase. Usage patterns might change from a mix of read and write operations to predominantly read-only operations, and access patterns might change from predominantly random requests to predominantly sequential requests. The Virtual Disk Copy feature allows a systems administrator to optimize virtual disk performance by moving a data set from one set of physical disks to another set. In a single operation, the administrator could optimize application performance and capacity use for virtually every aspect of virtual disk geometry for the data set—such as size, RAID level, number of drives, and so forth.

Application Development and Data Analysis

Typical uses for a virtual disk copy include decision support systems where analysis must be done on the source data without affecting the source's performance or application development where multiple full copies of the data can be subjected to different test scenarios with full performance in each.

CREATING A VIRTUAL DISK COPY

You can use either the Dell Modular Disk Storage Manager graphical user interface or the command line interface (CLI) to create a virtual disk copy. The CLI has the added benefit of command scripting, allowing for task automation and scheduling.

NOTE: Before creating a virtual disk copy, stop all I/O activity to the source and target virtual disks. You should also run the SMrepassist utility on the source virtual disk. SMrepassist (replication assistance) is a host-based utility for Windows that is designed to ensure that all the memory-resident data for the file system is flushed to disk.

When you configure a virtual disk copy, it is a best practice to use different physical disks and disk groups for the source virtual disk and the target virtual disk. By doing so, you can increase the availability of your data. If the disk group that contains the source virtual disk or the target virtual disk becomes unavailable you can still have access to the other disk group and your data.

Refer to the appropriate Dell™ PowerVault™ documentation for procedures on your chosen method for creating a virtual disk copy.

RE-COPYING A VIRTUAL DISK COPY

The Re-copy option enables you to create a new virtual disk copy for a previously defined copy pair that might have previously had a status of Stopped, Failed, or Completed. Use the Re-copy option for creating scheduled, complete backups of the target virtual disk that can then be copied to a tape drive for off-site storage. The Re-copy option is extremely beneficial in environments where backups or manipulation of the same data set frequently occur.

CAUTION: The Re-copy option overwrites existing data on the target virtual disk and makes the target *read-only* to hosts. The Re-copy option is designed to fail all snapshot virtual disks associated with the target, if any exist.

SECTION 5

COMBINING THE VIRTUAL DISK COPY AND VIRTUAL DISK SNAPSHOT PREMIUM FEATURES

CREATING A VIRTUAL DISK COPY FROM A SNAPSHOT

As mentioned earlier, the Virtual Disk Copy operation suspends I/O to the source virtual disk during a virtual disk copy to help prevent chronological inconsistencies from being created on the target virtual disk. In a 24x7 production environment, it is not always possible to suspend I/O to the source virtual disk for the amount of time required to create a virtual disk copy. To help alleviate this contention of resources, Virtual Disk Copy can be used in conjunction with Virtual Disk Snapshot.

It is a recommended best practice for all Virtual Disk Copy operations on production virtual disks to first create a snapshot of the virtual disk, and use the snapshot as the source of the Virtual Disk Copy operation. This combination can enable a complete point-in-time clone to be created with no detectable interruption to the I/O activity of the production virtual disk.

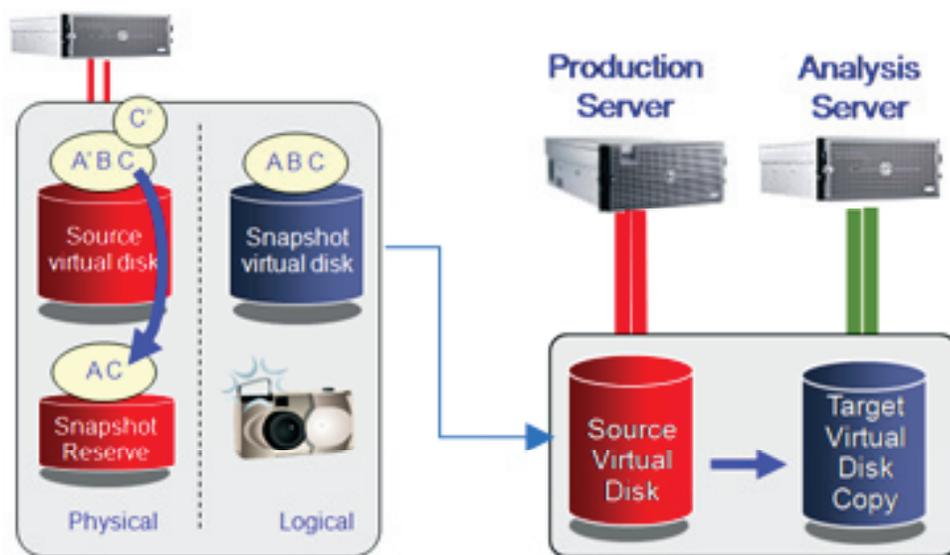


Figure 5: A Snapshot Virtual Disk as the Source for Creating a Virtual Disk Copy

Figure 6 shows the Dell Modular Disk Storage Manager view of a snapshot being selected as the source for a Virtual Disk Copy operation.

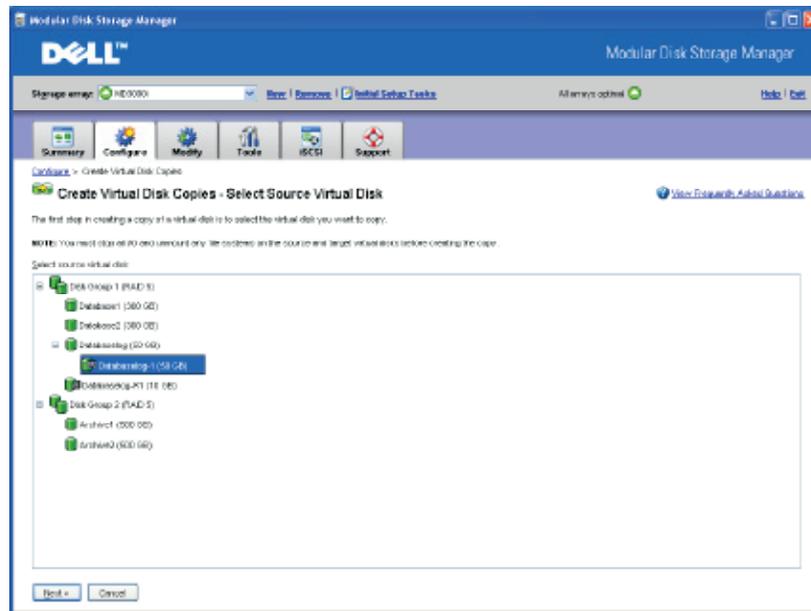


Figure 6: Selecting the Source Virtual Disk

RESTORING SNAPSHOT DATA TO THE SOURCE VIRTUAL DISK USING VIRTUAL DISK COPY

The Virtual Disk Copy premium feature can enable you to restore a virtual disk to the state represented by its snapshot. This can be used in the case where the data on a virtual disk is accidentally deleted or corrupted after a point-in-time snapshot is created. The data from the snapshot is now preferred to the data on the source virtual disk, and virtual disk copy can be used to restore the virtual disk to its original state using the snapshot.

Although Virtual Disk Copy can be leveraged to provide the desired functionality, the obvious method of simply copying from the snapshot back to the source virtual disk is not viable. Such a scenario would write to every block of the source virtual disk, creating a severe impact on the repository which tracks differences between the source virtual disk and all of its snapshots, including the snapshot from which it is being copied. As a result, the copy likely would fail as the repository fills to capacity.

The solution is to first create a new virtual disk to serve as a temporary target for the virtual disk copy. Then use Virtual Disk Copy to copy the snapshot to this temporary target virtual disk. Upon completion, remove this copy relationship and then create a new Virtual Disk Copy with the temporary target virtual disk serving as the source virtual disk and the original source virtual disk serving as the target.

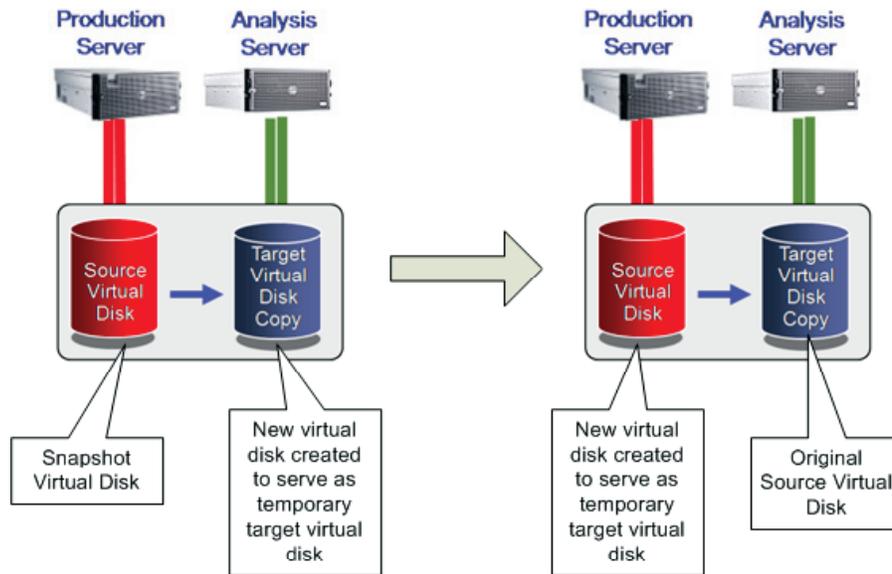


Figure 7: Virtual Disk Copy Process for Snapshot Virtual Disk to Source Virtual Disk

SECTION 6

CONCLUSION

The Dell Virtual Disk Snapshot and Virtual Disk Copy premium features offer many benefits for data protection and management.

BENEFITS OF VIRTUAL DISK SNAPSHOT

- Can help improve employee productivity by having an immediate point-in-time image of the data—no waiting for large virtual disks of data to copy.
- Helps protect data by providing a readily available online point-in-time image that is designed to reduce restore time, especially as compared to traditional tape media backups.
- Provides rapid application development through the immediate creation of data test sets.

BENEFITS OF VIRTUAL DISK COPY

- As a full copy (clone), a virtual disk copy can be assigned to any host and is intended to be used by applications requiring a copy of production data without affecting the performance of the production virtual disk.
- Virtual Disk Copy has five user-defined priority settings, and is controller-based, helping to minimize the overall I/O impact on the storage system and host applications.

BENEFITS OF COMBINING VIRTUAL DISK SNAPSHOT AND VIRTUAL DISK COPY

- You can use a snapshot to create a virtual disk copy without interrupting the I/O and response time of the production application.
- You can use Virtual Disk Copy to restore snapshot data to the original source virtual disk to recover deleted or corrupted data.

SECTION 7

CONTACT INFORMATION

Go to: <http://www.dell.com/>

or call 1-800-WWW-DELL

Information in this document is subject to change without notice.

SECTION 8

APPENDIX A: COMPARING A SNAPSHOT TO A VIRTUAL DISK COPIES

Feature	Virtual Disk Snapshot	Virtual Disk Copy
Replication type	Copy-on-write	Point-in-time fully copy (clone)
Copy location relative to the source	Same storage array	Same storage array
Performance impact	Processing against both source and pointer-based repository image might create some performance overhead	After the copy is complete, there should be no more performance overhead
Ongoing impact of maintaining copy	Some performance impact due to processing against both the data source and table of pointers	None
Max copies	4 per virtual disk / 128 per system	8 active copies per system / 255 non-active
Copy dependence on source	Lost if source is lost	Available if source is lost
Copy size	Fraction of the capacity of the source virtual disk (approx 20 percent or less of the source capacity in most cases)	Same amount of capacity as the source data required
Availability	Recovers source data in the event of corruption	Recovers source data in the event of corruption or hardware failure
Application access to source and target	Full read/write access while the snapshot is active	Source is read-only during copy process. No access to copy during copying process
Typical uses	<ul style="list-style-type: none"> • Point-in-time backup images / File Restoration • Operations where the data must be frozen at a particular point in time • Data mining and analysis 	<ul style="list-style-type: none"> • Full point-in-time data for analysis, mining and testing without directly affecting the source's data or performance • Migrating data between storage tiers • Full recovery protection against virtual disk corruption
Features	<ul style="list-style-type: none"> • Almost instantaneous – when limited time is available for a full virtual disk copy • Less disk space than a virtual disk copy • Can be mapped to any host • Can be read from or written to 	<ul style="list-style-type: none"> • Eliminates I/O contention on the primary virtual disk • Can be mapped to any host • Can be read from or written to

Table 1: Comparison of Virtual Disk Snapshot and Virtual Disk Copy

SECTION 9

REFERENCES

Dell, 2007. Dell PowerVault Modular Disk Storage Manager CLI Guide,
<http://support.dell.com/support/edocs/systems/md3000/en/index.htm>

Dell, 2007. Dell PowerVault Modular Disk Storage Manager User's Guide,
<http://support.dell.com/support/edocs/systems/md3000/en/UG/HTML/iscsi.htm>

Microsoft, 2003. Storage Management Using Windows Server 2003 and Windows Storage Server 2003
Virtual Disk Service and Volume Shadow Copy Service,
<http://www.microsoft.com/windowsserversystem/wss2003/techinfo/plandeploy/stormgtusingvdsvss.mspx>

SIMPLIFY STORAGE AT DELL.COM/MD3000i

