



Technical Overview of NetApp SnapDrive®

Network Appliance, Inc.

April 2007 | TR-3197

Abstract

SnapDrive provides storage virtualization for NetApp storage volumes and Snapshot™ backup and restore operations for hosts running various operating systems, including Windows®, Solaris™, HP-UX, AIX, Red Hat, and SUSE Linux® over iSCSI and Fibre Channel (FCP) transport protocols. This paper describes the architecture and presents a technical overview of the SnapDrive solution for Windows and UNIX® environments.

Contents

1. Purpose	3
2. Existing Scenario.....	3
3. NetApp Solution.....	3
4. SnapDrive Integration with Host Operating System.....	5
5. Storage Management Using SnapDrive.....	6
6. Snapshot Management Using SnapDrive	6
7. Complementary Solutions	9
8. Summary	10
9. Revision History.....	11

1. Purpose

This paper introduces NetApp SnapDrive software and highlights the benefits that the product brings to a host administrator for efficient storage and data management, including the ability to implement quick backup and restore policies for application data. The paper also briefly discusses the various functionalities and deployments of SnapDrive.

2. Existing Scenario

Today's enterprise IT administrator is challenged to provision storage quickly to support applications for new business initiatives with minimal budget. For most IT administrators, it doesn't stop at this point; they need to protect the application against data corruption, disasters, and attacks with the help of well-planned backup mechanisms. Even more importantly, the backup process should not disrupt the service or the performance of the service.

To address these issues, administrators maintain their own set of scripts that are executed periodically or manually to speed up and automate the tasks of storage provisioning and backing up. But this comes with an additional burden of maintaining the scripts. For example, if the host operating system is updated to a newer version, the output of the command used in the script or command syntax might change, resulting in a complete rewrite of the scripts.

Therefore IT and host administrators require a solution that enables them to:

- Quickly provision storage and bring it online from the host
- Adapt to different operating systems and other environment changes without requiring maintenance of scripts
- Take backups without any performance degradation to the application
- Schedule and maintain their own backup policies, depending on the type of application
- Modify space allocations without taking applications offline

3. NetApp Solution

SnapDrive Overview

NetApp provides SnapDrive for Windows and SnapDrive for UNIX solutions to address all of the issues discussed in the previous section. SnapDrive helps host administrators to provision storage and manage it directly from the host. SnapDrive gives flexibility to application administrators by enabling them to define their backup policies; and more importantly, it allows administrators to resize the storage on the fly without any disruption of application service. Both versions of SnapDrive simplify storage and data management by using the host operating system and NetApp technologies; by hiding the complexity of steps that must be executed on both the storage system and the host system; and by removing the dependency on the storage administrator.

Key SnapDrive functionality includes storage provisioning on the host, consistent data Snapshot copies, and rapid application data recovery from Snapshot copies. SnapDrive complements the native file system and volume manager technology, and integrates seamlessly with the clustering technology supported by the host operating system to provide high availability of the service to its users.

Note: The term SnapDrive as used in this document refers to both SnapDrive for Windows and SnapDrive for UNIX.




Figure 1) Typical SnapDrive deployment.

SnapDrive provides a layer of abstraction between an application running on the host operating system and the underlying NetApp storage systems. Applications that are running on a server with SnapDrive use virtual disks (or LUNs) on NetApp storage systems, as if they were locally connected drives or mount points. This allows applications that require locally attached storage, such as Oracle® database, Microsoft® Exchange, and Microsoft SQL to leverage NetApp technologies, including Snapshot, flexible volumes, cloning, and space management technologies.

SnapDrive includes all the necessary drivers and software to manage interfaces, protocols, storage, and Snapshot copies. Snapshot copies are nondisruptive to applications and functions on execution. Snapshot backups can also be mirrored across LAN or WAN links for centralized archiving and disaster recovery.

NetApp SnapManager® solutions are designed to invisibly leverage SnapDrive to trigger backups, restores, and mirroring of specific data sets. For more information, see <http://www.netapp.com/products/enterprise-software/manageability-software/app-suite.html>.

Benefits of SnapDrive

Most of today's enterprises use business-critical applications, and their IT and storage management teams face a number of challenges. They must:

- Support new business initiatives with minimal increase in operating budget
- Protect data from corruption, disasters, and attacks
- Back up data without any performance degradation, quickly and consistently, without any errors

SnapDrive addresses these problems by providing simplified and intuitive storage management and data protection from a host/server perspective. The following list highlights some of the important benefits of SnapDrive:

- Allows host and application administrators to quickly create virtual disks within a dynamic pool of storage that can be reallocated, scaled, and enlarged in real time, even while systems are accessing data
- Dynamic on-the-fly file system expansion—new disks are usable within seconds
- Snapshot copies provide rapid backup and restore capability with minimal resource and capacity requirements
- Mirroring, data replication, and clustering for high availability
- Supports multipath technology for high performance
- Enables connecting to existing Snapshot copies from the original host or a different host
- Independent of the underlying storage access media and protocol; SnapDrive supports FCP, iSCSI, and NFS as the transport protocols (NFS supports only Snapshot management.)
- Patented, high-performance, low-latency file system with industry-leading reliability
- Robust yet easy-to-use data and storage management features and software
- Industry-leading availability, exceeding 99.99% availability on nonclustered systems
- Robust data integrity features such as advanced RAID functionality and built-in file system checksums help protect against disk drive failures and disk errors

4. SnapDrive Integration with Host Operating System

Once installed, SnapDrive can be used to create and manage virtual disks on NetApp storage systems from the host operating system. For Windows-based hosts, these appear as basic disks to the Windows 2000 Server, the Windows Server 2003, and their applications. The virtual disks that reside on the NetApp storage system can be expanded, unlike Windows-native basic disks. SnapDrive is also used to create, delete, and manage all aspects of the application Snapshot backups.

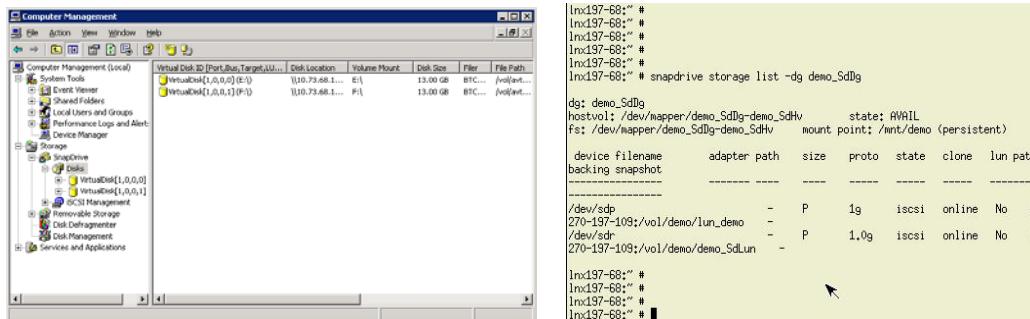


Figure 2) SnapDrive user interface in Windows and UNIX environments.

On UNIX platforms, the SnapDrive functionality is supported mostly through CLI commands, and you can provision storage directly on the host. You can manage the entire storage hierarchy, from the host-side application-visible file, down through the volume manager, to the storage-system-side LUNs that form the actual storage repository.

5. Storage Management Using SnapDrive

SnapDrive provisions storage to hosts by automating the task of LUN creation, LUN mapping to host initiators, and then mounting the LUN on the host. SnapDrive hides the complexity of storage administration tasks by automating all of these tasks, including LUN creation, LUN mapping, creation of disk groups and host volumes, and mounting the disk groups and LUNS to the file system on the local host.

SnapDrive brings dramatic simplicity and flexibility to host administrators in ways that are not available through alternative methods of administrating storage.

Storage provisioning: With SnapDrive, server administrators can quickly create and assign storage as the application needs grow.

Storage resizing: With SnapDrive, administrators can dynamically increase the size of the LUN at any time, even when the application is using the LUN data.

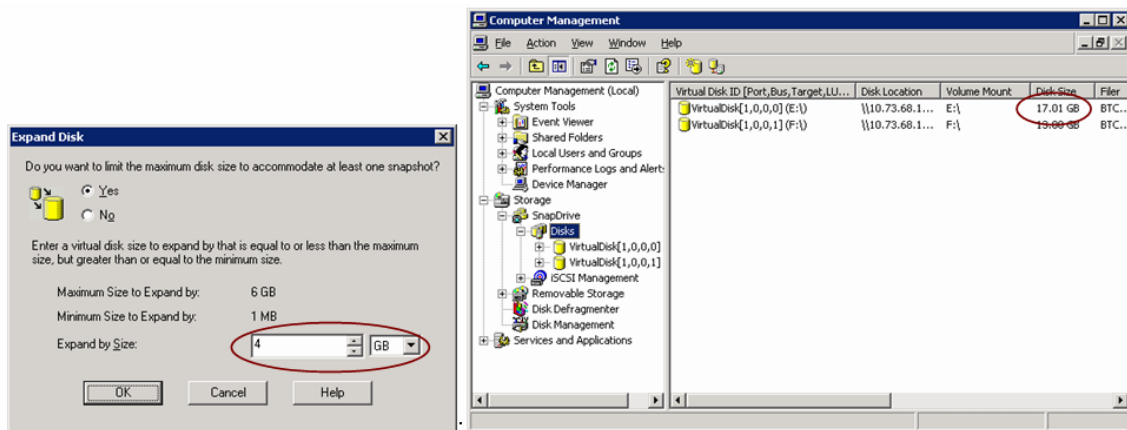


Figure 3) SnapDrive for Windows showing resizing virtual disk operation.

Migration: With SnapDrive, the administrator can disconnect the LUN from the host at any time, and connect that same LUN to a different host. This scenario is mostly seen when the administrator has a different setup with the upgraded software installed and wants to migrate the LUN to the new host without extensive downtime. In this case, the administrator simply disconnects the LUN from the old host and reconnects to the LUN from the new host using SnapDrive.

6. Snapshot Management Using SnapDrive

A thorough discussion of Snapshot technology is beyond the scope of this paper.

For the purposes of this document, Snapshot copies are not copies of data. Rather, they record the state of the blocks in the file system at a given point in time and provide read-only access to that image of the file system.

SnapDrive simplifies the integration and use of NetApp Snapshot technology, bringing a new level of data reliability and availability to the host applications. SnapDrive enables the user to create Snapshot copies of the file system, restore copies, and delete copies quickly and easily.

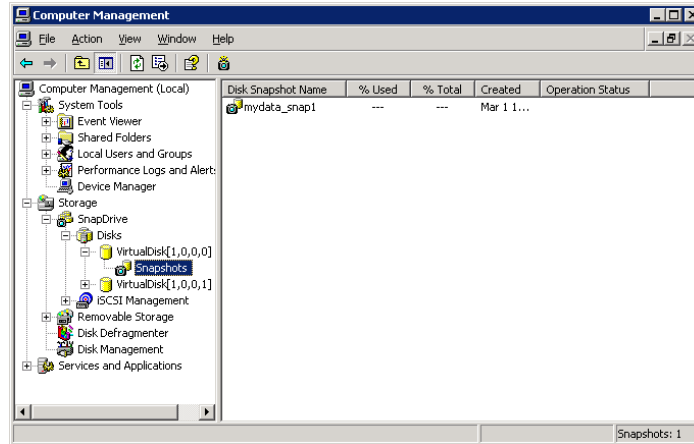


Figure 4) SnapDrive for Windows showing a Snapshot listing of a virtual disk.

Compared to Snapshot copies taken from the storage system, Snapshot copies created using SnapDrive are more useful in scenarios where the application data is tied to LUNs that have been created on different volumes, or even on different storage systems. SnapDrive takes special steps to ensure that Snapshot copies that span multiple storage systems, or storage system volumes, are also host-consistent. SnapDrive ensures this consistency by synchronizing the data in the cache and freezing I/O operations to the requested LUNs.

Note: Snapshot copies of entities contained in a single storage system volume are always crash-consistent. Creating Snapshot copies manually by using the storage system command line or FilerView® interface does not take the application into account and can produce unpredictable results.

Restoring Data from a Snapshot Copy

Snapshot copies are read-only and cannot be modified by browsing through the file system. SnapDrive includes the ability to restore data from a Snapshot copy nearly instantaneously. In addition, SnapDrive can create a writable version of a Snapshot copy from an existing Snapshot copy. These writeable copies are clones of the application data that can be used for a variety of purposes, such as test, development, and support. Using SnapDrive to recover data from Snapshot copies saves enormous expenses in the form of lost productivity and missed business opportunities.

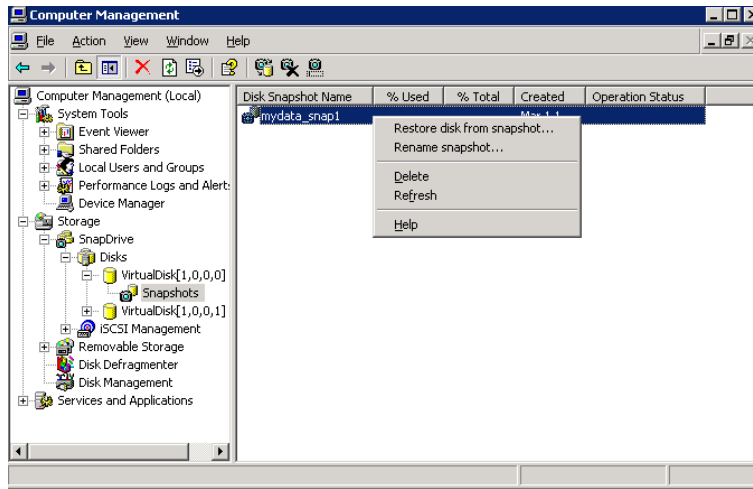


Figure 5) SnapDrive for Windows; restoring data from a Snapshot copy.

Cloning

Snapshot technology forms the basis for cloning. Cloning is used mostly in environments where a volume has to be duplicated for development, testing, and support, so that the testing procedure doesn't damage production data.

Writeable Snapshot software-backed LUNs enable the administrator to quickly clone application data, such as a database. They do this by taking a Snapshot copy, creating a Snapshot image-backed LUN, and then connecting that LUN to a host. Because of the structure of the WAFL[®] file system, the clone actually shares unchanged blocks with the original application data, but no data copy is involved. Changed blocks are written to new areas on disk, and the WAFL inode pointers are updated with this location information.

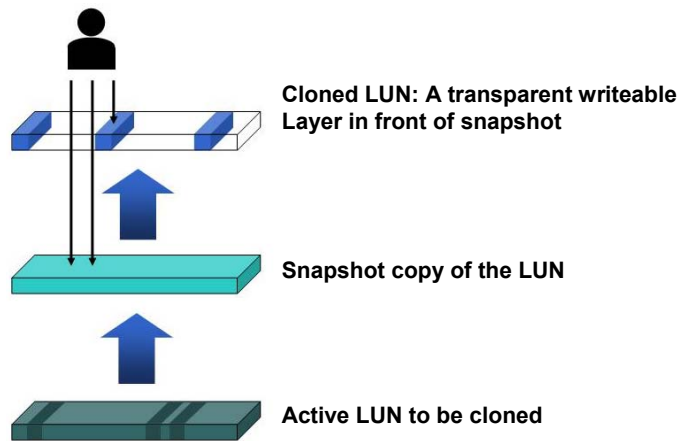


Figure 6) LUN clone.

With SnapDrive, the administrator can connect to a specific Snapshot copy by mounting it to a new host or file system and then restoring files or directories from that copy. Once the desired files or directories are restored, the Snapshot copy can be disconnected from the host.

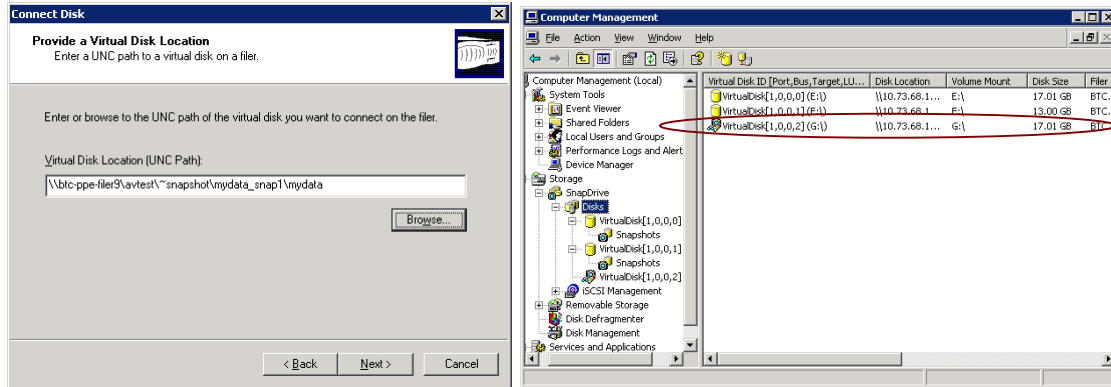


Figure 7) SnapDrive for Windows showing creation and listing of a cloned virtual disk.

Figure 7 shows how to create a clone of a virtual disk, by connecting to the LUN in the Snapshot directory. It also shows how the cloned virtual disk is listed in the SnapDrive for Windows application.

Usage of Clone

Snapshot database backups can be quickly restored and mounted for read and write access in minutes, so that administrators can re leverage the data for other purposes, such as running reports. That data can be accessed without affecting the production database.

Most enterprises engage in rigorous testing, where they need to duplicate the application data and test the new or modified application with the existing data. This means that when they use archaic methods to create copies, 100GB of application data becomes 200GB as they create a new copy. This can be quite costly. To keep the costs low, one undesirable option would be to reduce the number of clones or to reduce the scope of the data being cloned. To avoid this scenario, NetApp cloning technology helps to replicate the full application data, and the administrator can create and delete clones without compromising performance and storage space.

7. Complementary Solutions

SnapDrive plays a critical role in NetApp SAN deployments and makes the NetApp storage and Snapshot technology application aware. This section describes how other NetApp applications use SnapDrive to provide a complete solution for applications.

SnapManager

SnapManager for Oracle (SMO), SnapManager for SQL Server (SMSQL), and SnapManager for Exchange (SME) are part of the NetApp Application Suite. They enable administrators to simplify data management tasks, speed up application-consistent restore and recovery, and easily add or expand capacity, greatly increasing their productivity and enhancing data availability.

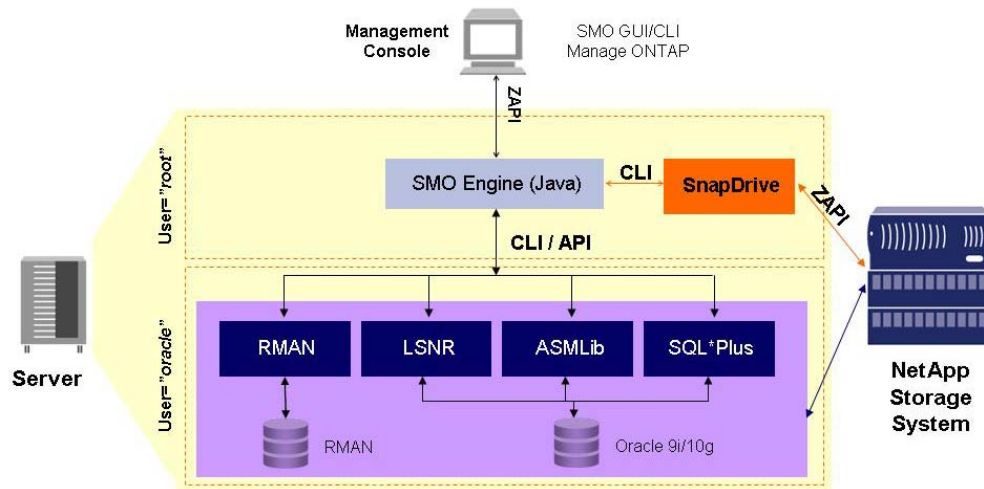


Figure 8) SnapDrive deployment with SnapManager for Oracle.

Role of SnapDrive for UNIX

SnapManager for Oracle depends on SnapDrive for UNIX to provide the underlying functionality of storage provisioning and Snapshot management. For example, SnapManager identifies backup data sets and puts the Oracle database in hot backup mode while an application-consistent Snapshot copy is created by SnapDrive.

The following operations are among those executed from SnapManager for Oracle that are executed by SnapDrive for UNIX:

1. Request Snapshot copy creation
2. Connect host to Snapshot copies
3. Clone from Snapshot copies
4. Restore Snapshot copies
5. Disconnect from Snapshot copies

8. Summary

SnapDrive integrates host-based data and file systems with NetApp storage systems by using virtual disks. This makes all of the related storage appear as if resides on the host server via the drives and mount points.

Administrators can execute nearly instantaneous Snapshot backups and restorations of application data. Snapshot backups can be mirrored to one or more locations across a LAN or WAN link for centralized archiving or disaster recovery. SnapDrive coordinates Snapshot execution with supported applications.

9. Revision History

Date	Name	Description
February, 2007	Arun Raman	Update
September, 2002	John Phillips	Creation

