Back to Basics: SnapProtect

Chris Blackwood
Technical Marketing Engineer

This article is the seventh installment of Back to Basics, a series of articles that discuss the fundamentals of popular NetApp® technologies.

An important reason why IT teams choose NetApp storage is because it provides the ability to use integrated data protection technologies like Snapshot™ copies, SnapMirror® replication, and SnapVault® disk-to-disk backup. These capabilities dramatically accelerate and simplify backup and replication for DR and other purposes.

Still, we saw a need for deeper integration with backup applications, especially for those who need to include tape in their backup environments.

NetApp introduced its SnapProtect® management software about a year ago to provide these and other features. NetApp partnered with CommVault to integrate core components of CommVault® Simpana® with core NetApp technologies. The combination delivers all the benefits you expect from Snapshot copies, SnapMirror, and SnapVault, plus it offers significant advantages including:

- Accelerates both backup and restore operations
- Full tape support
- Cataloging of Snapshot copies, replicas, and tape
- Built-in support for VMware®, Hyper-V®, and popular applications
- Automated provisioning of secondary storage
- Cascading and fan-out configurations
- Flexible scheduling and retention
- Reporting
- Simple single-pane-of-glass management for all features

This chapter of Back to Basics explores how NetApp SnapProtect is implemented, common use cases, best practices for implementing SnapProtect, and more.

How SnapProtect Is Implemented

SnapProtect uses a variety of components. Most of these are familiar NetApp technologies such as:

- Snapshot copies

Explore

More Back to Basics

Learn the fundamentals of NetApp core technologies by reading other chapters in this series:

- Chapter 1: Thin Provisioning
- Chapter 2: Deduplication
- Chapter 3: FlexClone
- Chapter 4: Volume SnapMirror
- Chapter 5: RAID-DP
- Chapter 6: Data Compression
- Chapter 7: SnapProtect

NetApp Data Protection Technology: Partner Integration

In addition to working with CommVault on SnapProtect, NetApp works with several additional backup partners to make advanced NetApp data protection features accessible from their products:

- Symantec NetBackup Replication Director—Integrates NetApp Snapshot and replication technologies into the Symantec™ NetBackup™ framework.
- NetApp Syncsort Integrated Backup—Leverages NetApp SnapVault block-incremental replication to offer D2D and D2D2T backup for heterogeneous (non NetApp) data.

- SnapMirror replication
- SnapVault for disk-to-disk backup
- FlexClone® technology for cloning and indexing
- SnapRestore® technology for rapid restore of whole volumes and single files
- OnCommand® software (formerly NetApp Operations Manager) for provisioning and replication

In addition, SnapProtect adds several additional components that enable cataloging, coordination, management, and so on.

- SnapProtect Server. Runs Windows®, Microsoft® SQL Server®, and management software
- MediaAgents. Additional servers that help spread the data protection workload during SnapProtect operations
- iDataAgents (iDAs). Software agents installed on backup clients that are responsible for data consistency during backup operations

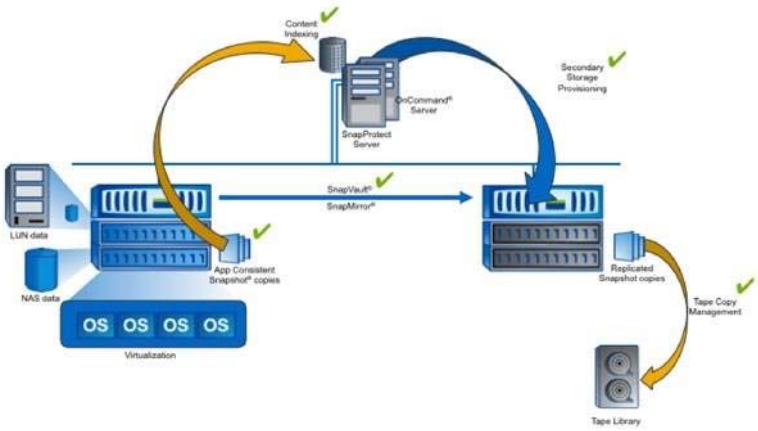

Figure 1) Overview of SnapProtect operations.

Most SnapProtect operations begin with the creation of a Snapshot copy on primary storage. This happens very quickly and provides the first level of data protection. Each Snapshot copy is initiated by SnapProtect according to established schedules. Once a Snapshot copy completes it is indexed in the SnapProtect database.

Next, if a SnapVault or SnapMirror activity is going to occur, SnapProtect passes control to the NetApp OnCommand server to provision necessary secondary storage using the Resource Pools and provisioning policies assigned to SnapProtect. OnCommand then builds datasets, provisions necessary volumes, and starts baseline transfers from primary to secondary.

If a tape copy is needed, SnapProtect provides three options, two of which can also be indexed in the SnapProtect catalog. (See Table 1.)

Table 1) Tape options.

Option	iDataAgent	Indexed?	Notes
NDMP dump	NetApp NAS NDMP iDA	Yes	
SMTape (SnapMirror to Tape)	NetApp NAS NDMP iDA	No	Data ONTAP® 8.0.1 software and later
Streaming through media agent	All other iDAs	Yes	

SnapProtect support for various data types

The specifics of SnapProtect operation vary slightly depending on the type of data being protected.

NAS data. When a Snapshot copy is made of NAS data, SnapProtect simply indexes the contents of the Snapshot copy directly using the SnapDiff API, which returns a list of files that have changed (been created, modified, or deleted) since the last Snapshot copy was made.

For NAS data you can restore directories, files, and qtrees as needed, or you can revert an entire volume using SnapRestore. If a file to be restored is available in a Snapshot copy on primary storage, Single File SnapRestore (SFSR) is used; a copy-back operation is used for restores from secondary storage.

LUN data. Performing a SnapProtect operation on LUN data requires an iDA on the host accessing the LUN. This is because the host understands the format of the data being written. For example, the Windows File System iDA understands NTFS.

The file system is quiesced (using Volume Shadow Copy Service on Windows hosts) before the Snapshot copy is made. LUN clones are created and used for indexing the contents inside the LUN.

Application data. Each application has its own iDA on the application host. The iDA prepares the database for backup prior to the creation of a NetApp Snapshot copy. Restore options vary depending on the application being protected. Supported applications are shown in Table 2. Note that the backup and restore of application data by SnapProtect does NOT use NetApp SnapManager® products.

Table 2) Supported applications.

Application	Support
Microsoft Exchange	Windows
Microsoft SQL Server	Windows
Microsoft SharePoint®	Windows
DB2	UNIX® and Linux®
Oracle® (including RAC)	UNIX and Linux
SAP® on Oracle	UNIX and Linux

Virtual machine data. A key feature of SnapProtect is the ability to protect many virtual machines quickly without the need to install an agent in each virtual machine. In addition, SnapProtect can index the contents of each VM with different levels of recoverability, including single file recovery. You can learn more about the advantages of SnapProtect in virtual environments in a recent NetApp white paper.

SnapProtect software uses a virtual server agent (VSA), which runs on a media agent designated to protect the virtual environment, to perform data protection operations for virtual environments. Within the VSA, instances are created that define the type of virtualization solution being used.

Discovery rules can be established so that new virtual machines are automatically added and protected. For example, using a discovery rule of “Datastore Affinity” automatically protects new virtual machines on specific datastores.

Use Cases

Of the various possible use cases, NAS, virtual environments, and application backup and restore are probably the ideal uses for SnapProtect. The tool not only gives you full cataloging, but you also get a simple interface for recoveries with wild-card search capabilities that let you identify and restore necessary files in less time. Because tape management and replication are all integrated, you can manage most of your data protection needs from a single pane of glass.

VMware. For VMware backup and restore, SnapProtect communicates with vSphere® via the VSA, and VMware-level snapshots are created before NetApp Snapshot copies. Single file recovery and live-browse are supported for Windows VMs, allowing individual files from within a VM to be restored. Single File SnapRestore (SFSR) can be used for VMDKs stored in NFS datastores to accelerate the restore of an entire virtual machine.

Application backups using an application agent (iDA) as described in the previous section can be performed, but only for applications not installed on VMDKs. Exchange and SQL Server can be quiesced via the VMware VSA when they are installed in VMDKs, providing consistent backup.

Hyper-V. VSS is required to back up an online Windows VM. (Windows 2000 must be offline.) Single file restore is supported for Windows and Linux VMs.

Cascade and Fan-Out

In addition to its ability to protect particular types of data, another important SnapProtect use case is support for cascade and fan-out configurations that let you more easily achieve your data protection, disaster recovery, and compliance objectives.

These configuration options allow you to fully automate the creation of backup copies and/or mirrors as required. For example, you can mirror to a remote site and then vault at that location (or to a third location) to provide the backup history you need. You can also vault locally and then mirror the vault to a remote location, giving you full local and remote copies of your important backups. The options are almost endless. You can add backup to tape at any point in a cascade or fan out (primary, secondary, local, remote) so you also have tape copies where you need them.

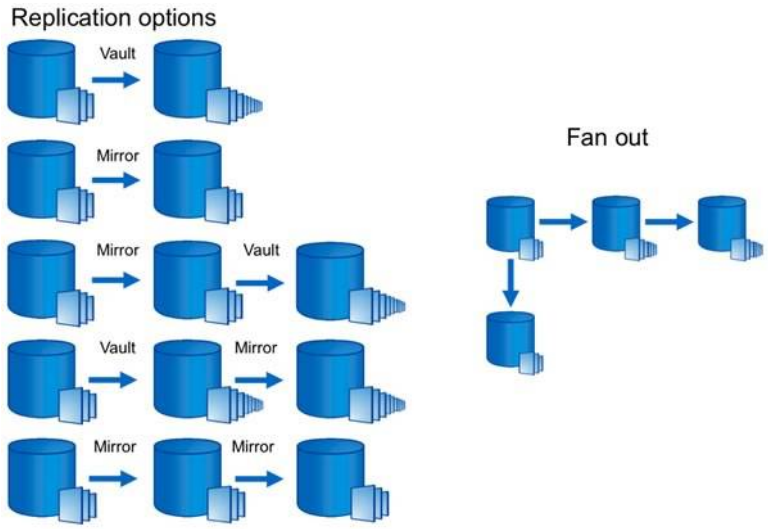

Figure 2) SnapProtect cascade and fan-out support.

Using SnapProtect

To get started using SnapProtect, your primary storage must be running Data ONTAP 7.2.6 or later or Data ONTAP 8.0.1 or later. You need only license the software for primary storage systems. However, if you have a MetroCluster™ configuration (that is supported) you will need to license SnapProtect on both nodes. Similarly, if you use SnapMirror to mirror to a DR site and you want to be able to run SnapProtect at the DR site after a failover, you should license SnapProtect there as well.

In addition to the SnapProtect license, you will need valid licenses for the other products utilized by SnapProtect, and these will need to be licensed on both primary and secondary storage. This includes SnapMirror and/or SnapVault, FlexClone, and SnapRestore. The OnCommand server will need to be licensed as well.

Scheduling and Retention

Flexible scheduling and retention policies are another strength of SnapProtect. Scheduling can be done by creating individual schedules or by creating schedule policies that group various schedules together. In addition to backups it is also possible to schedule other operations such as restores and client installs.

There are two types of retention rules in SnapProtect. Basic retention rules apply to all backups. Extended retention rules apply to longer-term retention, such as weekly full, monthly full, and yearly full backups, and they are intended to give you the flexibility to meet a wide range of business requirements. These rules are configured in the storage policy and can be set separately for the primary Snapshot copy, vault copies, and tape copies. Mirror copies do not allow specific retention settings, because they inherit the same retention as the primary copy.

For more details on getting started using SnapProtect, check out NetApp Technical Report 3920: NetApp SnapProtect Management Software: Overview and Design Considerations. You can also follow the discussion and ask questions on the SnapProtect thread on NetApp Communities.

SnapProtect and Other NetApp Technologies

SnapProtect works with most other NetApp technologies without changes. This includes common storage efficiency technologies such as deduplication, thin provisioning, and compression.

However, SnapProtect cannot coordinate all of these technologies directly. For instance, you can enable and manage secondary deduplication from SnapProtect, but you cannot do the same with compression. That would need to be managed separately.

Cluster-Mode. SnapProtect support for storage systems running in Cluster-Mode is limited only to tape support at this time.

Qtree SnapMirror. SnapProtect uses asynchronous volume SnapMirror for replication, but not qtree SnapMirror. If you wish to use qtree SnapMirror, it must be managed separately.

Conclusion

NetApp SnapProtect software is an important data protection tool that provides catalog and tape support so you have a single management framework to create and manage local Snapshot copies, replication (with SnapVault or SnapMirror) to secondary storage, and movement to tape. SnapProtect is ideal for IT teams that have a group responsible for backups across multiple environments (applications, physical, virtual).

To learn more about NetApp SnapProtect, refer to TR-3920: NetApp SnapProtect Management Software: Overview and Design Considerations. To find out more about the use of SnapProtect specifically for virtual environments, check out WP-7131: Accelerate Backup and Recovery in Virtual Server Environments with NetApp SnapProtect.

By Chris Blackwood, Technical Marketing Engineer

Chris has 18 years of experience in the IT industry. Since joining NetApp more than four years ago, he has focused on data protection solutions including SnapVault and Open Systems SnapVault. Now he dedicates his time exclusively to all aspects of SnapProtect, including interfacing with CommVault engineering and product management.

Quick Links

- Tech OnTap Community
- Archive
- User Groups
- PDF
