Technical Report

NetApp SnapVault with IBM Lotus Domino

John B. Spinks, NetApp
November 2010 | TR-3883

USING SNAPVAULT IN YOUR DOMINO ENVIRONMENT

NetApp® SnapVault® is a heterogeneous disk-to-disk backup solution for NetApp storage systems and heterogeneous operating systems (Microsoft® Windows®, Linux®, and UNIX®). SnapVault leverages NetApp Snapshot® copies and block-level incrementals for reliable, low-overhead backup and recovery suitable for any environment. It is a reliable and economical way to protect enterprise data, and it has many significant advantages over traditional backup methods. Use SnapVault in your Domino environment to store NetApp Snapshot copies on a remote NetApp storage system, which enables you to use less expensive disks for long-term storage.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

1 INTRODUCTION

1.1 DOCUMENT PURPOSE AND SCOPE

NetApp SnapVault technology is designed to safeguard data on NetApp storage faster and more cost effectively. SnapVault is an extremely reliable way to protect enterprise data, and it has many significant advantages over traditional backup methods. Although SnapVault can be deployed in configurations designed to emulate the legacy backup methods it replaces, you can realize the full value of the solution by making a significant shift in the way you think about backup and recovery.

This document describes the steps required to configure NetApp SnapVault on your NetApp storage system and provides information about the expected behavior of SnapVault in your IBM Lotus Domino environment.

1.2 ASSUMPTIONS

This document describes the implementation of NetApp SnapVault in an environment that is already in production. We assume that your NetApp storage system and IBM Lotus Domino environment are already configured and in working condition. Furthermore, we assume that you are familiar with the basic operations of a NetApp storage system, including NetApp Data ONTAP® fundamentals, the Domino server, and your host operating system.

2 SNAPVAULT OVERVIEW

SnapVault is a disk-based storage backup feature of Data ONTAP that enables data stored on multiple storage systems to be backed up to a central, secondary storage system (SnapVault Secondary) as read-only Snapshot copies. The SnapVault Secondary is a data storage system running Data ONTAP.

Initially, a complete copy of the data from the primary storage is backed up to the SnapVault Secondary, and a Snapshot image of the data volume is created. For the subsequent backups, SnapVault transfers only the data blocks that have changed since the previous Snapshot backup.

3 SNAPVAULT BENEFITS

SnapVault offers many benefits to NetApp customers. Historically, enterprise data backup has been a time-consuming, expensive, and sometimes unreliable process. Today, the volume of data being generated is rapidly increasing, while the time available to back up this data is decreasing. With SnapVault technology, customers can overcome this critical business challenge. SnapVault provides a centralized, disk-based backup solution for heterogeneous storage environments. Storing backup data in multiple Snapshot copies on the SnapVault Secondary storage lets you keep days, weeks, or even months of backup data online and ready for fast restoration. SnapVault gives you the power to choose which data gets backed up, the frequency of the backup, and how long backup copies are retained. Additional benefits of SnapVault are:

- Increase your confidence in data backup and recovery
- Lower your total cost of backups
- Reduce the impact of backups on your production environment
- Use policy-driven and centralized management
- Use centralized backup
- Enjoy easy use and administration

4 SNAPVAULT ARCHITECTURE

There are several ways to configure SnapVault in your Domino environment. Some of the most common configurations are:

- SnapVault Primary connected to SnapVault Secondary
- SnapVault Primary connected to multiple SnapVault Secondary systems
- SnapVault Primary and SnapVault Secondary with NetApp SnapMirror®
- SnapVault Primary and SnapVault Secondary with a tape drive
- SnapVault Primary and SnapVault Secondary with active-active high-availability (HA) pair storage

4.1 SNAPVAULT PRIMARY CONNECTED TO SNAPVAULT SECONDARY

Figure 1 illustrates the basic SnapVault configuration, in which SnapVault is used to back up data on the SnapVault Primary to the SnapVault Secondary. Businesses use this very common scenario to create a backup copy on disk rather than performing a backup to tape. The SnapVault Secondary can be located at the same physical data center or at a remote or disaster recovery site.

Figure 1) SnapVault Primary connected to SnapVault Secondary.

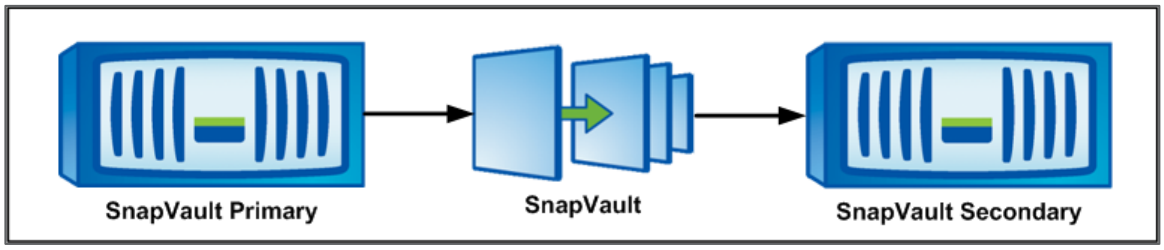

4.2 SNAPVAULT PRIMARY CONNECTED TO MULTIPLE SNAPVAULT SECONDARY SYSTEMS

Figure 2 illustrates the configuration in which one SnapVault Primary is connected to two or more SnapVault Secondary systems. This scenario is commonly used when backup data needs to be stored at multiple locations or when scalability is a concern.

Figure 2) SnapVault Primary connected to multiple SnapVault Secondary systems.

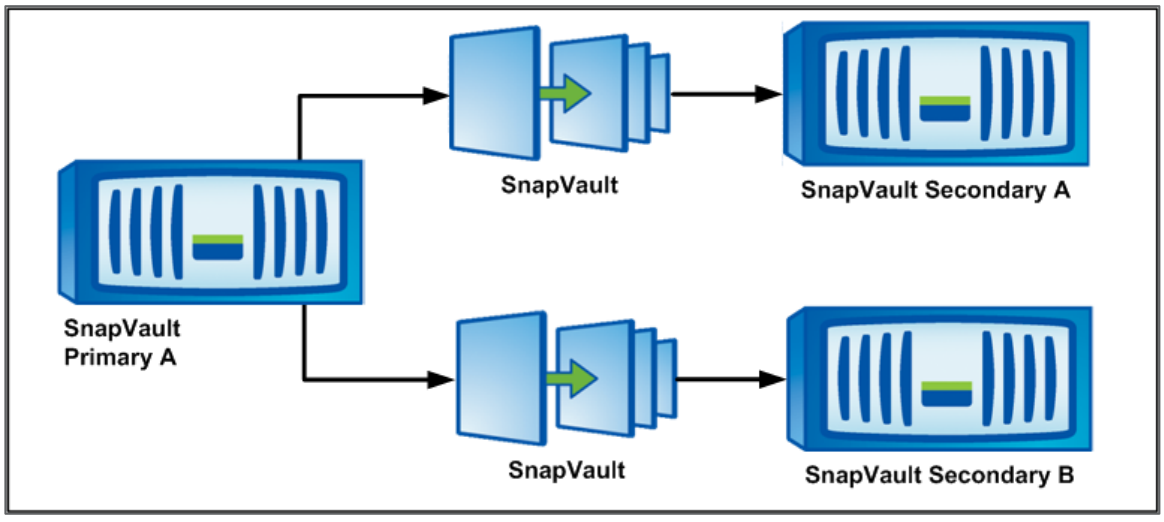

4.3 SNAPVAULT PRIMARY AND SNAPVAULT SECONDARY WITH SNAPMIRROR

This configuration offers an additional level of data protection. The SnapVault Primary backs up data to the SnapVault Secondary, and Snapshot backup images are replicated from the SnapVault Secondary to a SnapMirror destination storage system. SnapMirror replication is performed on a volume level. If the SnapVault Secondary stops functioning, data can still be accessed from the SnapMirror destination system. The SnapMirror system can also be converted to a SnapVault Secondary for further updates from the SnapVault Primary. Figure 3 illustrates the primary devices backing up data to a centralized SnapVault device, which is then mirrored to another location by using SnapMirror. This configuration is commonly used when centralized backup and redundancy are required.

Figure 3) SnapVault Primary and SnapVault Secondary with SnapMirror.

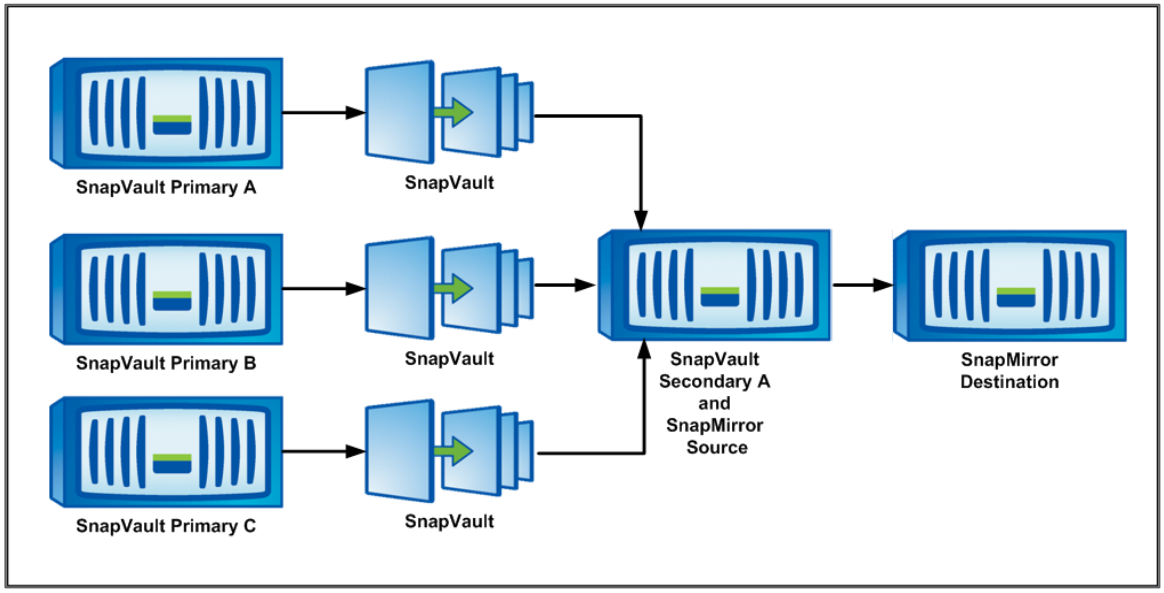

4.4 SNAPVAULT PRIMARY AND SNAPVAULT SECONDARY WITH A TAPE DRIVE

Figure 4 illustrates the configuration that offloads data from the SnapVault Secondary to a tape drive. With this configuration, less frequently used data can be moved from costly disks to low-cost tape, while recent data is retained in the SnapVault Secondary system. This configuration can assist in providing a truly centralized tape backup solution.

Figure 4) SnapVault Primary and SnapVault Secondary with a tape drive.

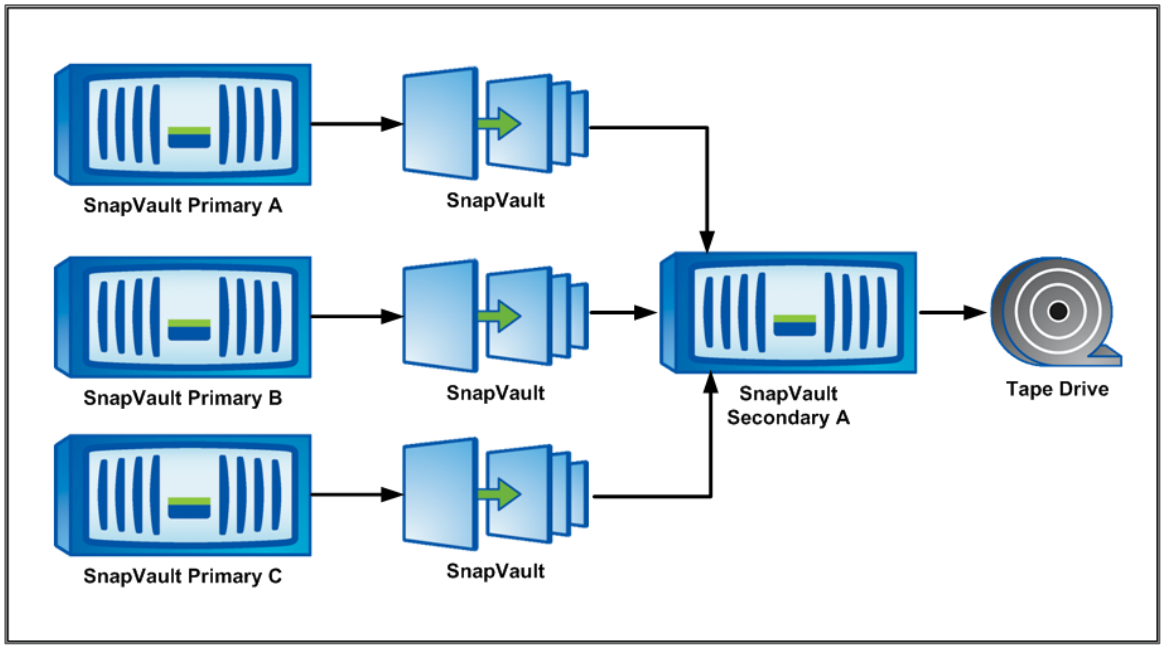

4.5 SNAPVAULT PRIMARY AND SNAPVAULT SECONDARY WITH ACTIVE-ACTIVE HA PAIR STORAGE

Figure 5 illustrates the configuration in which Node 1 of an HA pair is configured with the SnapVault Primary, while Node 2 is configured with the SnapVault Secondary. If an HA failover occurs, either single storage Node 1 or Node 2 can manage both the primary and secondary relationships, achieving data backup. Additionally, if Node 1 has Fibre Channel (FC) drives, you can move the data to SATA drives on Node 2, enabling you to easily use less expensive storage.

Figure 5) SnapVault Primary and SnapVault Secondary with active-active HA pair storage.

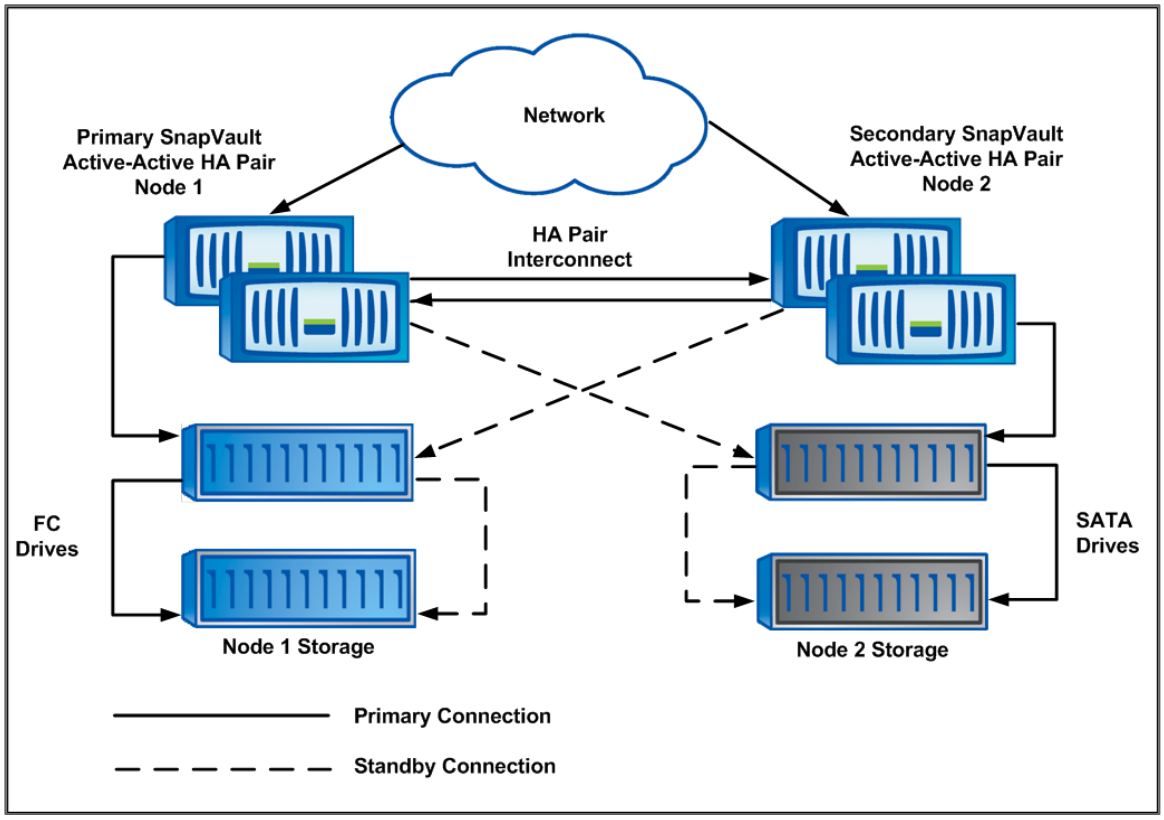

4.6 SNAPVAULT QTREE, VOLUME, AND LUN LAYOUT

A qtree is the lowest granularity container for SnapVault recovery. More than one qtree can exist per volume in the SnapVault relationship. Several logical configurations can be designed to back up SnapVault Primary Snapshot copies to SnapVault Secondary storage. These configurations are:

- Qtree-to-same-qtree backup
- Qtree-to-different-qtree backup
- Non-qtree data backup

QTREE-TO-SAME-QTREE BACKUP

As shown in Figure 6, you can place block-level logical unit numbers (LUNs), FC, or iSCSI inside qtrees for application data.

Figure 6 also shows that, in a qtree-to-same-qtree backup, a single volume with two qtrees on the SnapVault Primary backs up to the same two qtrees on the SnapVault Secondary.

This method potentially improves the recovery time objective, because the single qtree would be restored without affecting the other data on the volume. For example, you could restore a single LUN under qtree_c, as opposed to restoring the entire volume.

Figure 6) Qtree-to-same-qtree backup.

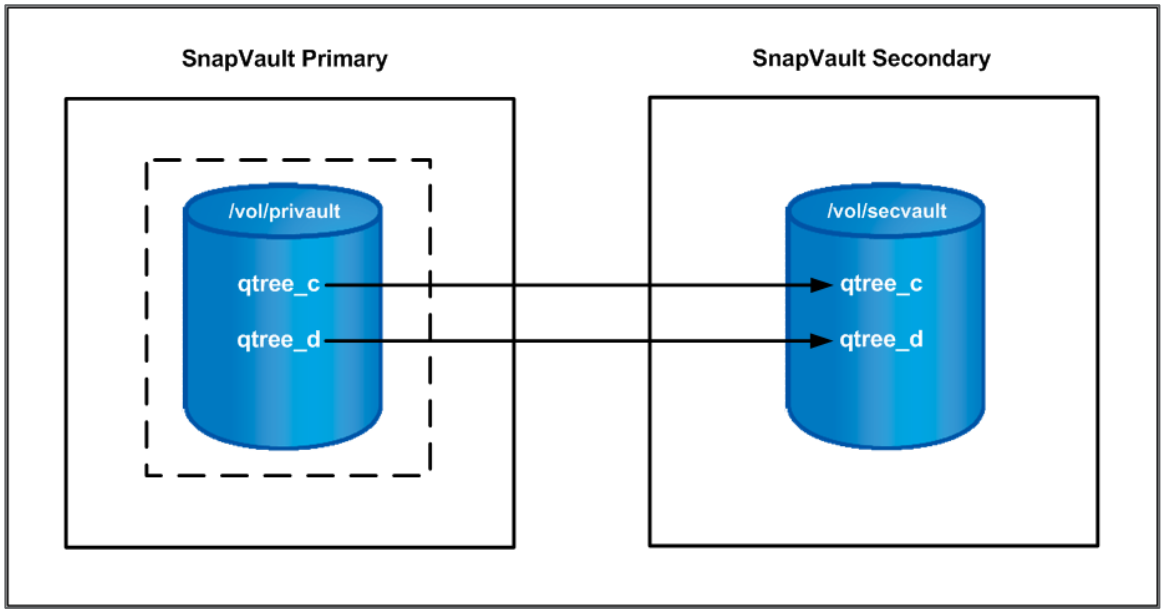

QTREE-TO-DIFFERENT-QTREE BACKUP

This configuration establishes different backup relationships between the primary and secondary qtrees. Administration is also facilitated by separating SnapVault Primary qtrees onto different volumes on the SnapVault Secondary volumes. As shown in Figure 7, the same volume on the SnapVault Primary is hosting Oracle® and Microsoft Exchange data, each in its own qtree. The SnapVault transfer moves each qtree backup into its own volume for the appropriate data type.

Figure 7) Qtree-to-different-qtree backup.

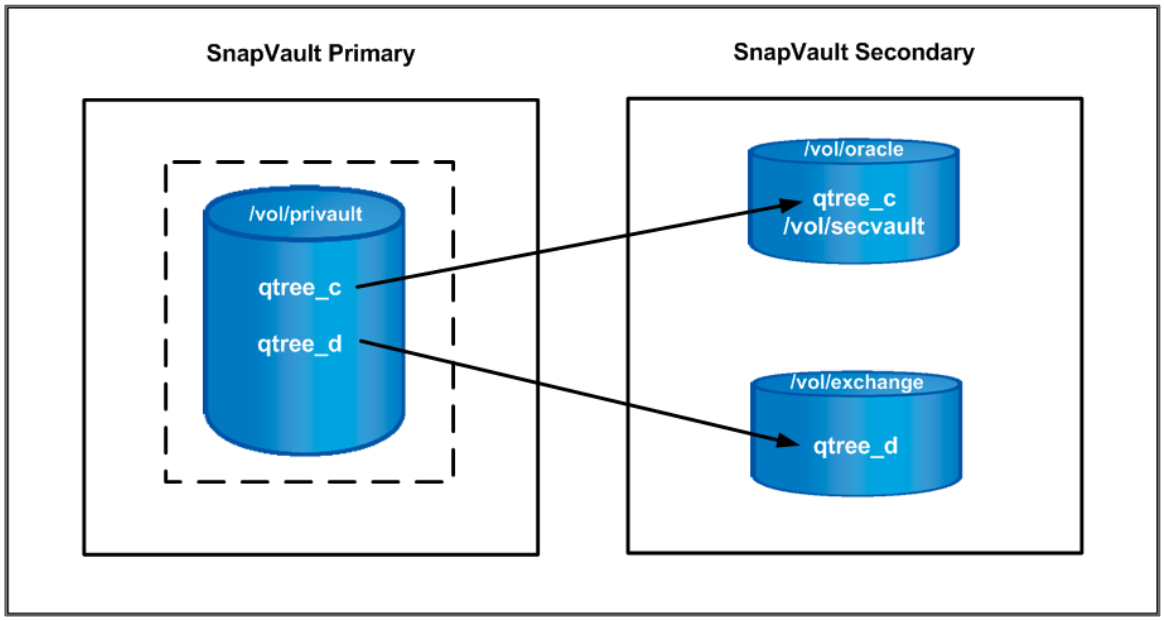

SNAPVAULT FOR NON-QTREE DATA

This configuration is most often used when you want to use SnapVault to back up an existing dataset, but the data is not in a qtree. You can specify an option while performing the data transfer in this configuration. This type of transfer is performed by using a dash (-) instead of a qtree name in the SnapVault start command. This command tells SnapVault to replicate all of the data in the root of the volume, but not in a qtree. A downside of this method is that you must restore to a qtree; therefore, the added functionality of an incremental restore is lost. As shown in Figure 8, the volume data is captured in a qtree on the SnapVault Secondary.

Figure 8) SnapVault for non-qtree data.

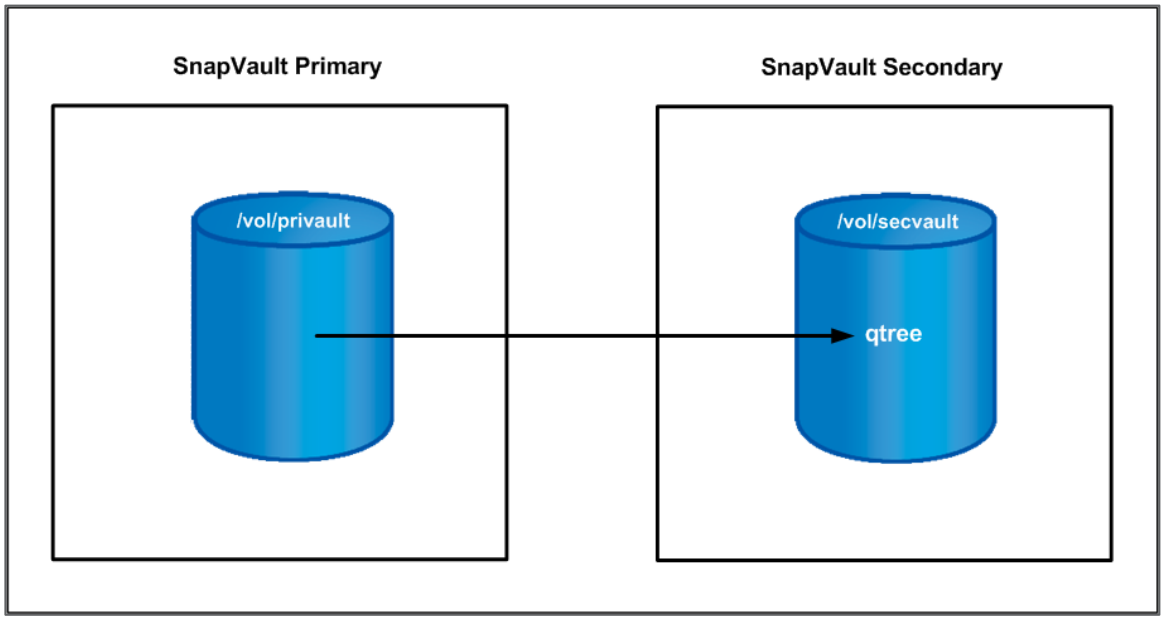

5 CONFIGURATION OVERVIEW

This section describes the initial configurations performed on the SnapVault Primary and SnapVault Secondary. We created a basic configuration by using SnapVault Primary to SnapVault Secondary with qtree-to-qtree backup.

5.1 SNAPVAULT STORAGE ARCHITECTURE

The diagram in Figure 9 and the information in Table 1 illustrate the example architecture used in this paper.

Figure 9) Environment overview.

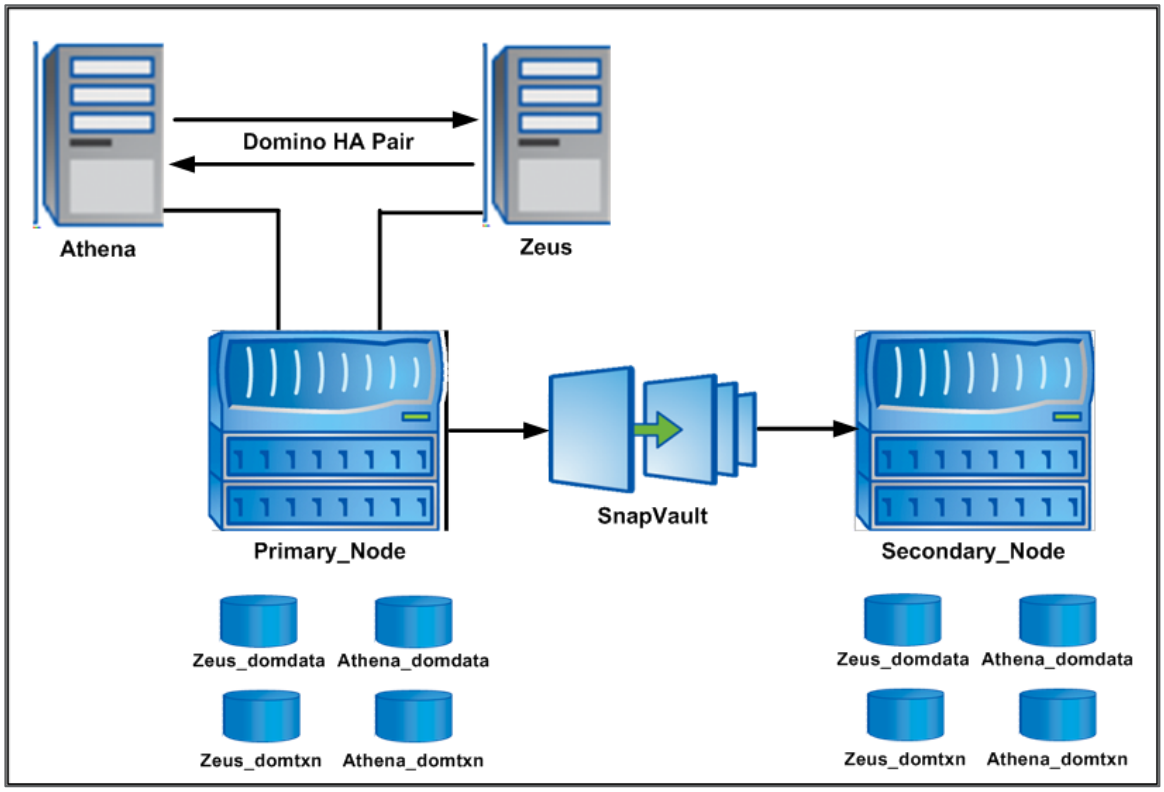

The example architecture consists of the following:

- **Two Domino servers.** Athena and Zeus.
- **Four volumes on Primary_Node (a SnapVault Primary).** /vol/Athena_domdata, /vol/Athena_domtxn, /vol/Zeus_domdata, and /vol/Zeus_domtxn.
- **One associated qtree on each volume.** For example, /vol/Athena_domdata has qtree Athena-domdata_qtree.
- **Four volumes on Secondary_Node (a SnapVault Secondary).** /vol/Athena_domdata, /vol/Athena_domtxn, /vol/Zeus_domdata, and /vol/Zeus_domtxn. These volumes hold the backup Snapshot copies from all qtrees defined on Primary_Node.

Each qtree on Secondary_Node is created by Data ONTAP during the initial `SnapVault start` command.

Table 1 provides the details of the environment.

Table 1) Environment overview.

Storage Name	Role	Volume	Qtree
Primary_Node	SnapVault Primary	/vol/Athena_domdata	Athena-domdata_qtree
		/vol/Athena_domtxn	Athena-domtxn_qtree
		/vol/Zeus_domdata	Zeus-domdata_qtree
		/vol/Zeus_domtxn	Zeus-domtxn_qtree

Storage Name	Role	Volume	Qtree
Secondary_Node	SnapVault Secondary	/vol/Athena_domdata	Athena-domdata_qtree
		/vol/Athena_domtxn	Athena-domtxn_qtree
		/vol/Zeus_domdata	Zeus-domdata_qtree
		/vol/Zeus_domtxn	Zeus-domtxn_qtree

5.2 INITIAL STORAGE CONFIGURATION

The storage systems must be able to communicate with the Domino database server and vice versa. A storage system can communicate with a database server if an entry exists in its `/etc/hosts` file for the database server or, alternatively, if it uses some other host name resolution techniques such as Network Information System (NIS) or Domain Name System (DNS). By default, the `/etc/hosts` file is checked first for host name resolution. The easiest way to update the `/etc/hosts` file on the storage system is by using NetApp System Manager.

Note: NetApp System Manager is a simple Windows-based tool that enables administrators to easily set up and manage NetApp storage systems. System Manager can be downloaded from the NetApp Support site (NOW®).

Entries made in the `/etc/hosts` file should be similar to the following:

```
[HostIP] [HostName]
```

Where

- `HostIP` identifies the IP address assigned to the database server.
- `HostName` identifies the name assigned to the database server.

For example, to add an entry for a storage system named Primary_Node with IP address 10.61.162.40, add the following line to the `/etc/hosts` file on the server or storage system:

```
10.61.162.40 Primary_Node
```

Repeat this step on the servers and storage systems to update all appropriate `/etc/hosts` files.

Note: Parameters shown in angle brackets (`< >`) are optional. Parameters and options shown in square brackets (`[ ]`) are required and must be provided. A comma followed by an ellipsis (…) indicates that the preceding parameter can be repeated multiple times.

SNAPVAULT LICENSE REQUIREMENT

To properly configure SnapVault, enable SnapVault options on both the primary and secondary storage systems.

SnapVault requires an sv_ontap_pri license for the primary node and an sv_ontap_sec license for the secondary node.

You can add the SnapVault license by executing the `license add` command on the SnapVault Primary and Secondary storage systems.

```
license add [licenseCode]
```

Where

- `[licenseCode]` is the license code to add to your NetApp storage system.

For example, to apply license code ABCDE1234, execute the following command on the storage system:

```
license add ABCDE1234
```

Apply the proper licenses to each of your storage systems.

After the licenses are added to each storage system, enable SnapVault by setting the `snapvault.enable` option to on. Execute the following command on both the primary and the secondary storage systems to enable SnapVault:

```
options snapvault.enable on
```

After SnapVault is enabled, set the `snapvault.access` option to control which systems can request data transfers. Execute the following command on the primary system to grant access only to the secondary system:

```
options snapvault.access host=[snapvault_secondary]
```

Where

- `[snapvault_secondary]` is the host name of the SnapVault Secondary system.

For example, execute the following command on Primary_Node to set Primary_Node to grant SnapVault permission to Secondary_Node:

```
options snapvault.access host=Secondary_Node
```

The snapvault.access option also needs to be set on the secondary system to grant the primary system access in case of a restore. Execute the following command to set the secondary system to grant access to the primary system:

```
options snapvault.access host =[snapvault_primary]
```

Where

- `[snapvault_primary]` is the host name of the SnapVault Primary system.

For example, execute the following command on Secondary_Node to set Secondary_Node to grant access to Primary_Node:

```
options snapvault.access host=Primary_Node
```

SNAPVAULT CONFIGURATION

When configuring SnapVault, plan carefully for the Snapshot and SnapVault scheduling. Using SnapVault and regular Snapshot schedules at the same time on the same volumes causes conflict, which results in a schedule failure. Also, if you are using NetApp deduplication, schedule your deduplication to run before SnapVault to get better storage efficiency.

SnapVault Primary Node Configuration Changes

To properly schedule SnapVault, turn off the default Snapshot schedules, which are replaced by SnapVault Snapshot schedules for all volumes in your Domino environment.

To change the Snapshot copy schedule for a specific volume, enter the following command:

```
snap sched [vol_name][weekly] [nightly] [hourly]@[n,n,…]
```

Where:

- `[vol_name]` is the name of the volume on which to set the schedule.
- `[weekly]` is the number of weekly Snapshot copies to keep.
- `[nightly]` is the number of nightly Snapshot copies to keep.

- `[hourly]` is the number of hourly Snapshot copies to keep.
- `[n,n,…]` specifies the hours at which to create the hourly Snapshot copies.

Note: A zero in any of the three schedules (weekly, nightly, or hourly) disables Snapshot copies for that interval.

For example, to disable the Snapshot schedule for the volume Athena_domdata, set a schedule of 0 0 0:

```
snap sched Athena_domdata 0 0 0
```

Repeat this command for each volume.

After the default Snapshot schedule is disabled, set a SnapVault Snapshot schedule. These Snapshot copies are kept on the primary storage system for quick access.

For our test scenario, we used the Snapshot copy retention on the primary node, as shown in Table 2.

Table 2) Primary_Node SnapVault schedule.

Schedule Name	Snapshot Schedule	Number of Snapshot Copies	Hours to Take Snapshot Copies	Days to Take Snapshot Copies
sv_hourly	Hourly	23	0–22	Monday–Sunday
sv_daily	Daily	7	23	Monday–Sunday

This schedule configures hourly and daily Snapshot copies for the SnapVault relationships. We kept an entire day's worth of hourly Snapshot copies, except for 11 p.m. (23), and a week's worth of daily Snapshot copies, taken at 11 p.m. (7). This keeps an entire day's worth of Snapshot copies easily accessible on the primary storage system.

Create the SnapVault schedule for each qtree or volume that is part of your SnapVault relationship. To create the SnapVault schedule for a qtree or volume, execute the following command on the SnapVault Primary storage system:

`SnapVault snap sched [vol_name]_[schedule_name][n]@[day_list]@[hour_list]`

Where:

- `[vol_name]` is the name of the volume on the primary storage system of which you are taking a Snapshot copy.
- `[schedule_name]` is the name of the Snapshot schedule; for example, sv_hourly. The name of this Snapshot copy must be the same on both the primary and secondary storage systems.
- `[n]` is the number of Snapshot copies to retain.
- `[day_list]` is a comma-separated list that specifies the days on which a Snapshot copy is created. Entries are not case sensitive, and a range can be specified by using a dash (-); for example, mon-fri. The dash (-) by itself indicates that Snapshot copies are not created automatically. The default value is mon-sun.
- `[hour_list]` specifies the hours when a Snapshot copy is created. Valid entries are numbers 0 to 23, and a range can be specified by using a dash (-). The default value is midnight (0).

For example, to create an hourly SnapVault schedule named sv_hourly (to take Snapshot copies every hour from midnight to 11 p.m., per Table 2) for volume Athena_domdata, execute the following command on the SnapVault primary node:

```
SnapVault snap sched Athena_domdata sv_hourly 23@0-22
```

Note: You can set up your SnapVault schedule in Protection Manager, but NetApp recommends using the CLI instead.

As shown in this example, a day listing was not entered; therefore, the default schedule of mon-sun is in place for the hourly Snapshot copies.

To create a daily SnapVault schedule, named sv_daily, that takes Snapshot copies at 11 p.m. and keeps seven Snapshot copies (per Table 2) for volume Athena_domdata, execute the following command on the SnapVault primary node:

```
SnapVault snap sched Athena_domdata sv_daily 7@23
```

Once again, a day listing was not entered, and the default schedule of mon-sun is used instead.

SnapVault Secondary Node Configuration Changes

By default, Data ONTAP reserves 20% of space for Snapshot copies from the total space allocated for the volume. This reserve space, called SnapReserve, can be used only by Snapshot copies and not by the active file system. SnapVault volumes retain more Snapshot copies for longer periods of time than a typical volume. The majority of space with SnapVault is used for Snapshot copies; therefore, the 20% reserve does not provide any benefit. For this reason, NetApp recommends turning off SnapReserve so that your storage administrator can see a clear picture of the storage utilization with SnapVault.

Set SnapReserve to 0% on the SnapVault Secondary node by using the following command:

```
snap reserve [vol_name] [percent]
```

Where

- `[vol_name]` is the name of the volume.
- `[percent]` is the percentage of space to be reserved for Snapshot copies.

For example, to set SnapReserve to 0% on volume Athena_domdata, execute the following command:

```
snap reserve Athena_domdata 0
```

Normal Snapshot schedules also need to be turned off, because they are replaced by SnapVault Snapshot copies. Turn off the Snapshot schedule by executing the following command on the SnapVault Secondary system:

```
snap sched [vol_name][weekly] [nightly] [hourly]@[n,n,…]
```

Where:

- `[vol_name]` is the name of the volume on which to set the schedule.
- `[weekly]` is the number of weekly Snapshot copies to keep.
- `[nightly]` is the number of nightly Snapshot copies to keep.
- `[hourly]` is the number of hourly Snapshot copies to keep.
- `[n,n,…]` specifies the hours at which to create the hourly Snapshot copies.

Note: A zero in any of the three schedules (weekly, nightly, or hourly) disables Snapshot copies for that interval.

For example, to disable the Snapshot schedule for the volume Athena_domdata, set a schedule of 0 0 0:

```
snap sched Athena_domdata 0 0 0
```

Repeat this command for each volume.

The next step is to set the SnapVault schedule for the secondary storage system.

For our test scenario, we used the Snapshot retention on the secondary node, as shown in Table 3.

Table 3) Secondary_Node SnapVault schedule.

Schedule Name	Snapshot Schedule	Number of Snapshot Copies	Hours to Take Snapshot Copy	Days to Take Snapshot Copy
sv_hourly	Hourly	4	0–22	Monday–Sunday
sv_daily	Daily	14	23	Sunday–Friday
sv_weekly	Weekly	26	23	Saturday

For this scenario, we kept only 4 hourly Snapshot copies, because 23 hourly Snapshot copies and 1 daily Snapshot copy are kept on the primary storage system. Keeping 4 hourly copies gives us the option of restoring recent copies in case an issue occurs with the primary storage system.

Schedule the backup of the SnapVault Snapshot copies from the primary storage system to the secondary storage system by executing the following command:

```
SnapVault snap sched –x [sec_vol][schedule_name][n]@[day_list]@[hour_list]
```

Where:

- `-x` is a required portion of the command on the secondary system. `–x` specifies that the SnapVault secondary qtrees on the specified volume are updated from their associated primary storage system.
- `[sec_vol]` is the name of the volume on the secondary storage system of which you are taking a Snapshot copy.
- `[schedule_name]` is the name of the Snapshot schedule; for example, sv_hourly. The name of this Snapshot copy must be the same on the primary and secondary systems.
- `[n]` is the number of Snapshot copies to retain.
- `[day_list]` is a comma-separated list that specifies the days on which a Snapshot copy is created. Entries are not case sensitive, and a range can be specified by using a dash (-); for example, mon-fri. The dash (-) by itself means that a Snapshot copy is not created automatically. The default value is mon-sun.
- `[hour_list]` specifies the hours when a Snapshot copy is created. Valid entries are numbers 0 to 23, and a range can be specified by using a dash (-). The default value is midnight (0).

For example, to schedule the SnapVault Snapshot copy for volume Athena_domdata to update from the primary storage system and to keep four copies on the sv_hourly schedule (from midnight to 10 p.m.), execute the following command:

```
SnapVault snap sched –x Athena_domdata sv_hourly 4@0-22
```

To schedule the SnapVault schedule for sv_daily, which is a daily transfer (except Saturday) taken at 11 p.m, which keeps the 14 most recent daily Snapshot copies, execute the following command:

```
SnapVault snap sched –x Athena_domdata sv_daily 14@23@sun-fri
```

To schedule the SnapVault schedule for sv_weekly, which is a weekly transfer taken at 11 p.m. on Saturday, which keeps 26 Snapshot copies, execute the following command:

```
SnapVault snap sched Athena_domdata sv_weekly 26@23@sat
```

All the data necessary to store a weekly Snapshot copy is already on the secondary system. Therefore, creating a weekly schedule on the primary system is not necessary to keep a weekly schedule on the secondary system.

INITIAL SNAPVAULT BASELINE TRANSFER

Up to this point, we have licensed and enabled SnapVault and configured the SnapVault schedules on both the primary and secondary systems. SnapVault does not yet know which qtrees to back up or where to store them on the secondary system.

To provide SnapVault with this information, initialize SnapVault by using the `SnapVault start` command on the SnapVault Secondary node:

```
SnapVault start –S [pri_system]:[pri_qtree] [sec_system]:[sec_qtree]
```

Where

- `-S` specifies the primary system and path. Set the `–S` option the first time you run the `SnapVault start` command for each primary system qtree that you want to copy.
- `[pri_system]` is the name of the primary storage system.
- `[pri_qtree]` is the qtree on the primary storage system that SnapVault backs up.
- `[sec_system]` is the name of the SnapVault secondary, or destination, system to which data is transferred. The local host name is used if a secondary system is not specified.
- `[sec_qtree]` is the destination qtree on the secondary storage system.

Note: The qtree specified for `[sec_qtree]` must not exist on the secondary storage system before you run the `SnapVault start` command.

For example, to execute the `SnapVault start` command for Athena_domdata, execute the following command:

```
SnapVault start –S Primary_Node:/vol/Athena_domdata /vol/Athena_domdata
```

6 DOMINO AND SNAPVAULT CONSIDERATIONS

When you take a Snapshot backup of a database or application, make sure that the Snapshot copy is application consistent. Many databases have a special hot backup or quiescing mode that keeps the Snapshot backup images consistent by temporarily suspending data transfer to storage while a Snapshot copy is taken. Domino offers application programming interfaces (APIs) that allow you to place Domino in hot backup mode, but very few solutions use these APIs because they are complex. If your backup solution uses the Domino APIs, then your Snapshot copies are application consistent. NetApp SnapManager® for Domino, discussed in section 7, can be used as a solution. Unfortunately, many customers do not have access to backup software that uses the Domino APIs. Without using APIs, customers have two choices to back up the Domino database:

- Back up the database while Domino is running.
- Shut down the Domino server and back up during the backup and maintenance window.

Snapshot copies can be taken while Domino is running. If a Snapshot copy is taken while the database is open, then Domino marks the backup copy of the database as being in an unknown status. As a result, when Domino accesses the backup image the next time, for example during a restore operation, a consistency check is required. This automatic consistency check means that Domino is proactively monitoring the stability of the Domino environment. If the consistency check completes successfully, the databases are ready for use. Any databases that require additional action are taken offline and listed on the Domino console. Basic Domino maintenance tasks such as `fixup` and `updall` need to run to bring inconsistent databases back online. For additional information, refer to the Administrator Guide for Domino Server Maintenance.

Alternatively, you can shut down the Domino server before you take the Snapshot copy and restart the server after the Snapshot copy is completed. The Snapshot copy takes only about 1 second to complete, regardless of the size of the data, and it can be scheduled during your maintenance window. This method

provides consistency because activity does not occur while the Domino server is stopped. Rather than using the normal Snapshot copy method described in section 5.2, use a script to shut down the server, to take the Snapshot copy, and to start the Domino server. SnapVault is a storage-level operation and does not require any Domino-level action. A script could run as a background process after a Snapshot copy is completed, or it could be scheduled by using a cron job or Windows task scheduler. Your corporate service-level agreements will probably determine which method to use.

6.1 DOMINO BACKUP WITH SNAPVAULT

If you have followed all of the steps in the paper up to this point, then your system is currently configured to do the following:

- **Primary_Node.** This system is configured to take 1 Snapshot copy each hour: 23 Snapshot copies are kept as hourly Snapshot copies and 1 is kept as a daily Snapshot copy. Primary_Node keeps one day's worth of hourly Snapshot copies (23) and one week's worth of daily Snapshot copies (7). SnapVault updates are transferred automatically to Secondary_Node.
- **Secondary_Node.** This system is configured to receive SnapVault copies from Primary_Node. The SnapVault schedule on Secondary_Node is configured to keep 4 hourly Snapshot copies, 14 daily Snapshot copies, and 26 weekly Snapshot copies.

No additional work needs to be performed to back up your Domino system.

If you are using SnapManager for Domino, follow the directions in section 7 to use SnapVault and SnapManager for Domino together.

6.2 DOMINO RESTORE FROM SNAPVAULT

To restore data from the secondary storage system back to the primary storage system, use the `SnapVault restore` command.

If both the primary and the secondary storage systems are running Data ONTAP 7.3 (or a later version), you have the option of performing an incremental restore that transfers only incremental changes from the secondary qtree back to the primary qtree.

If you intend to restore a qtree to the exact qtree location on the primary storage system from which you backed it up, perform one of the following steps:

- **Baseline restore.** The baseline restore can be to an existing qtree or to a nonexisting qtree.

Note: In the case of a baseline restore to an existing qtree, the restore operation overwrites the qtree data. If one or both storage systems are running a version of Data ONTAP earlier than 7.3, you must perform a baseline restore.

- **Incremental restore.** The restore operation transfers only incremental changes from the secondary qtree to the specified primary qtree.

Note: The incremental restore is more efficient. Therefore, NetApp recommends first attempting an incremental restore, before restoring over an existing primary qtree. If the incremental restore fails because of a lack of Snapshot copies, then attempt an in-place baseline restore.

To perform a SnapVault restore, execute the following command on the primary storage system:

```
SnapVault restore [-f] [-k n] [-r] [-w] [-s snapname] -S
[sec_system]:[sec_qtree] [pri_system]:[pri_qtree]
```

Where:

- `-S [sec_system]:[sec_qtree]` specifies the secondary storage system and qtree path from which to restore the data.
- The `-f` option forces the command to proceed without first asking for confirmation from the user.

- The `-k` option sets the maximum transfer rate in kilobytes per second [n].
- The `-r` option attempts an incremental restore. The incremental restore can be used to revert the changes made to a primary storage system qtree from the point of any backed-up version on the secondary storage system. This option is more efficient than a baseline restore because it transfers only the required changes from the secondary storage system to the primary storage system.
- The `-w` option causes the command not to return after the baseline transfer starts. Instead, it waits until the transfer completes (or fails). At that time, it prints the completion status and then returns.
- The `-s` option specifies that the restore operation must use the specified Snapshot (snap name) on the secondary storage system.
- `pri_system` is the name of the primary storage system to which you want to restore. If specified, this name must match the name of the host system.
- `pri_qtree` is the name of the primary storage system qtree to which you want to restore.

After performing the SnapVault restore, either resume or discontinue the SnapVault relationship.

To resume the SnapVault relationship between the restored qtree and its backup qtree on the secondary storage system, enter the following command:

```
SnapVault start -r -S [pri_system]:[pri_qtree] [sec_system]:[sec_qtree]
```

Where

- `-r` is required to restart backups from a restored primary storage system.
- `-S` specifies the primary system and path. The `–S` option must be set the first time the `SnapVault start` command is run for each primary system qtree that you want to copy.
- `[pri_system]` is the name of the primary storage system.
- `[pri_qtree]` is the qtree on the primary storage system that SnapVault backs up.
- `[sec_system]` is the name of the SnapVault Secondary, or destination, system to which data is transferred. If no secondary system is specified, use the local host name.
- `[sec_qtree]` is the destination qtree on the secondary storage system.

To discontinue the SnapVault relationship between the restored qtree and its backup qtree on the secondary storage system, enter the following command:

```
SnapVault release sec_qtree [pri_system]:[pri_qtree]
```

Where

- `[sec_qtree]` is the name of the qtree on the secondary storage system to be released from the SnapVault relationship.
- `[pri_system]` is the name of the primary storage system.
- `[pri_qtree]` is the qtree on the primary storage system to be released from SnapVault.

For example, to restore the SnapVault Snapshot sv_daily.2 back to the primary storage system, execute the following command:

```
SnapVault restore -r –s sv_daily.2 –S
Secondary_Node:/vol/Athena_domdata/Athena-domdata_qtree
Primary_Node:/vol/Athena_domdata/Athena-domdata_qtree
```

The complete output of the restore command with all of the dialog options is as follows:

```
Primary_Node> SnapVault restore -r -s sv_daily.2 –S
Secondary_Node:/vol/Athena_domdata/Athena-domdata_qtree
Primary_Node:/vol/Athena_domdata/Athena-domdata_qtree
```

```
Restore will overwrite existing data in /vol/Athena_domdata/Athena-
domdata_qtree

Are you sure you want to continue? yes

Wed Mar 17 10:50:21 EDT [Primary_Node: vdisk.qtreePreserveComplete:info]:
Qtree preserve is complete for /vol/Athena_domdata/Athena-domdata_qtree.

Wed Mar 17 10:50:22 EDT [Primary_Node:
replication.dst.resync.success:notice]: SnapVault resync of
/vol/Athena_domdata/Athena-domdata_qtree to
Secondary_Node:/vol/Athena_domdata/Athena-domdata_qtree was successful.

Transfer started.

Monitor progress with 'SnapVault status' or the snapmirror log.

Primary_Node> Wed Mar 17 10:50:33 EDT [Primary_Node:
vdisk.qtreeRestoreComplete:info]: Qtree restore is complete for
/vol/Fuji15Data1/qtree-fuji15data1.
```

Now that the qtree restore is complete, execute the following command to restart the SnapVault relationship:

```
SnapVault start –r –S Primary_Node:/vol/Athena_domdata/Athena-domdata_qtree
Secondary_Node:/vol/Athena_domdata/Athena-domdata_qtree
```

6.3 DOMINO SINGLE FILE RESTORE

The most common restore activity with Domino is not a volume restore; it's a restore of a single database or mailbox. Thanks to the database structure of Domino and the flexibility of NetApp Snapshot technology, it's easy to restore a Domino database from a Snapshot copy either in the primary storage system or from SnapVault.

Several methods can be used to access a Snapshot copy of a database. A quick and common method is to directly access the Snapshot copy. Depending on which protocol is in use, this can be done by browsing the directory structure or by using NetApp SnapDrive®. This method is simple if you are using a file-based protocol such as NFS or CIFS. Your mounted storage system contains a hidden folder called `.snapshot`. To directly access the Snapshot copy:

1. Browse to your mounted volume and then change directories to `.snapshot`. A directory listing displays all of the Snapshot copies on this volume.
2. Change directories to the Snapshot copy of your choice.
3. Browse the file system to locate the file you want to access.

Note: Remember that the Snapshot copy is read only. NetApp recommends copying the database to your local workstation to manipulate it.

Using a block-based file system requires SnapDrive to mount the Snapshot copy. The next section explains how to use SnapDrive for Windows to mount a Snapshot copy.

6.4 USING SNAPDRIVE TO RESTORE A SINGLE DOMINO DATABASE

It's easy to restore a single Domino database, or even a single e-mail, by using the built-in functionality of SnapDrive.

Note: We used SnapDrive for Windows for the purpose of this paper. However, the same functionality is available in SnapDrive for UNIX.

In Domino, each mail file or application is its own database (with the .nsf extension). You can use SnapDrive to browse the Snapshot directory and mount the Snapshot copy as a read-only file system. You can then copy a Domino database from the Snapshot copy onto a local workstation or even into the production file system, possibly overwriting a corrupt database with a Snapshot copy. To restore a single Domino database, follow these steps:

1. Open Computer Management (Start > Settings > Control Panel > Administrative Tools > Computer Management).
2. Under Storage, expand SnapDrive and then expand Disks.

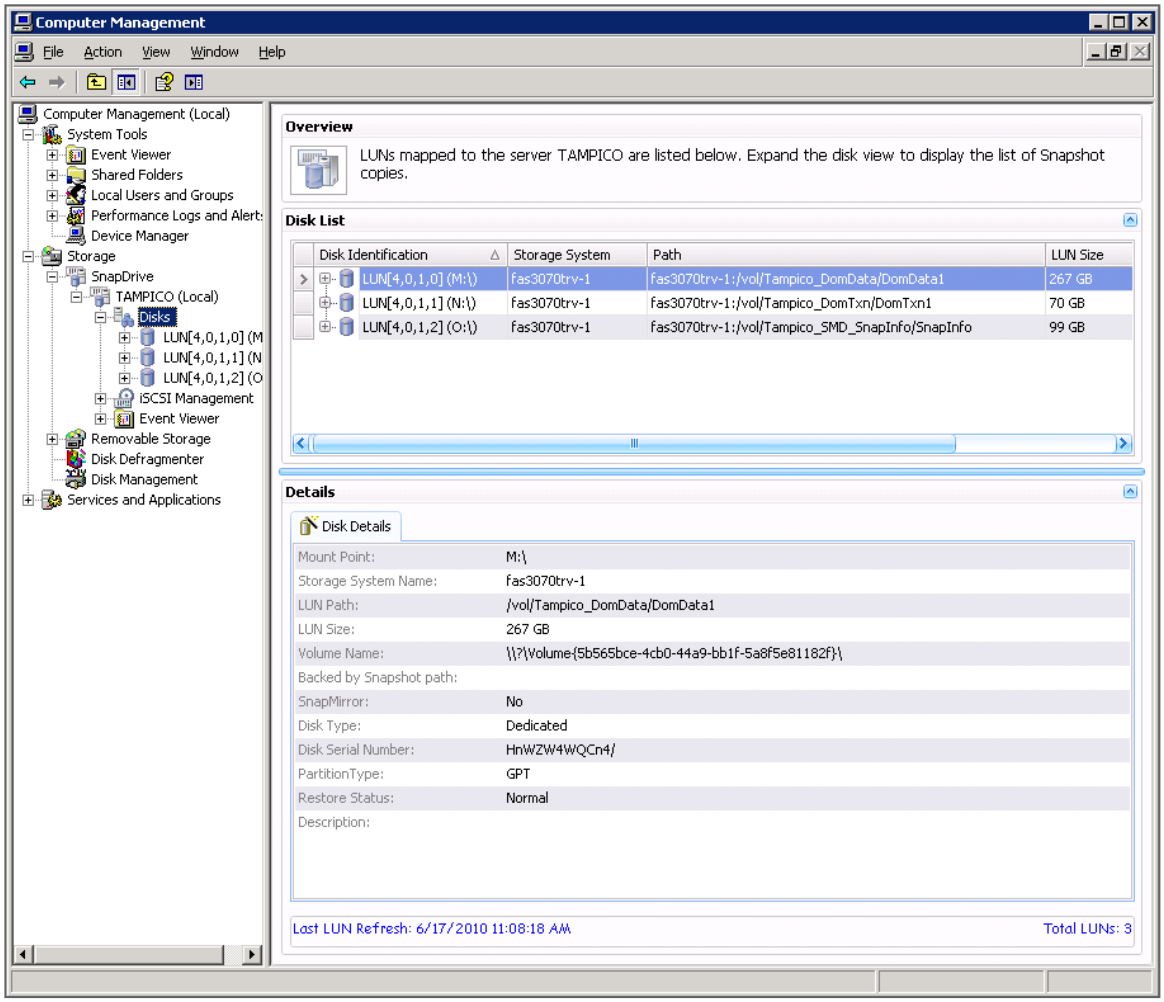

3. Right-click Disks and select Connect Disk.

The Connect Disk Wizard opens.

4. Click Next.

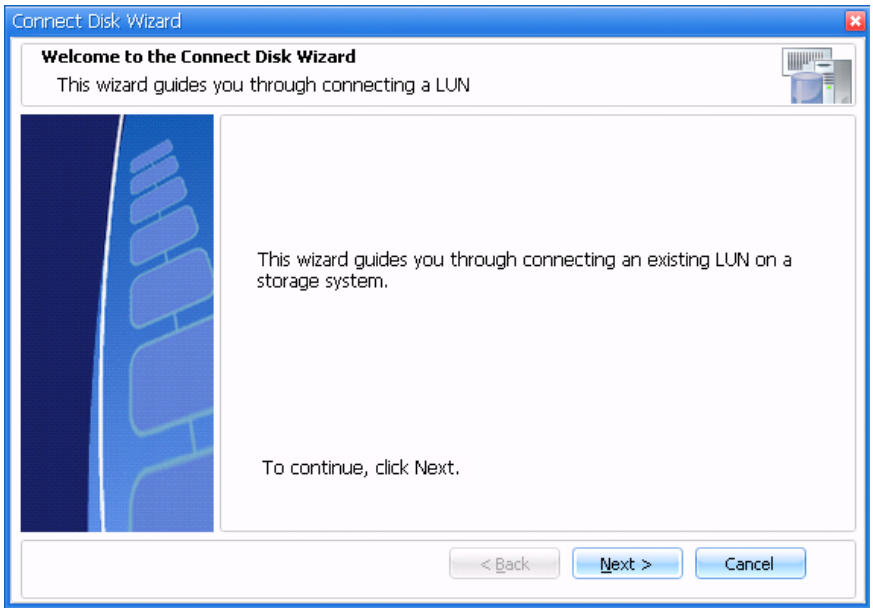

5. Select your storage system from the drop-down list.
6. Select and expand your volume.
7. Expand the `.snapshot` directory.
8. Expand the Snapshot copy to recover.
9. Select the LUN to mount.
10. Click Next.

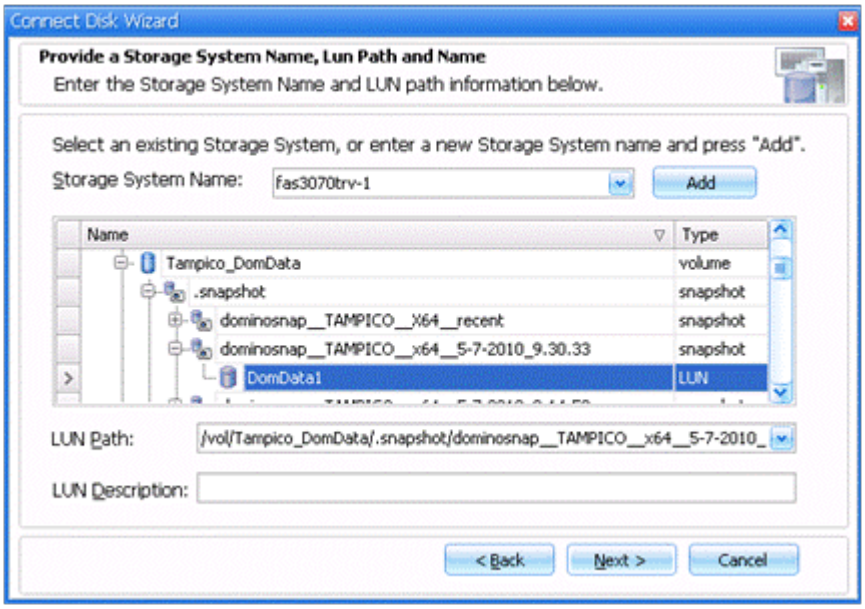

11. Select your virtual disk type. In this case, select Dedicated (your only option for a Snapshot copy).
12. Click Next.

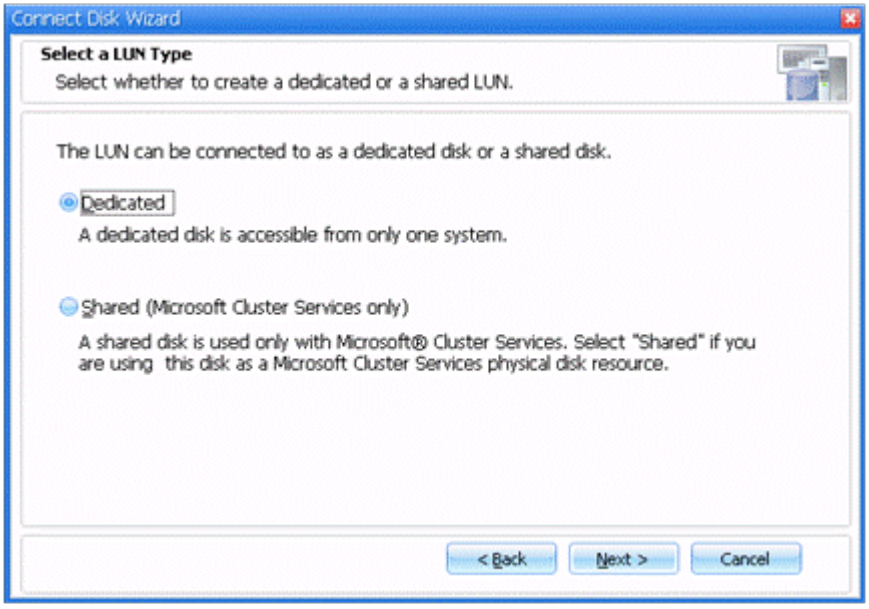

13. Select a virtual disk drive letter, in this case Z, for easy identification.
14. Click Next.

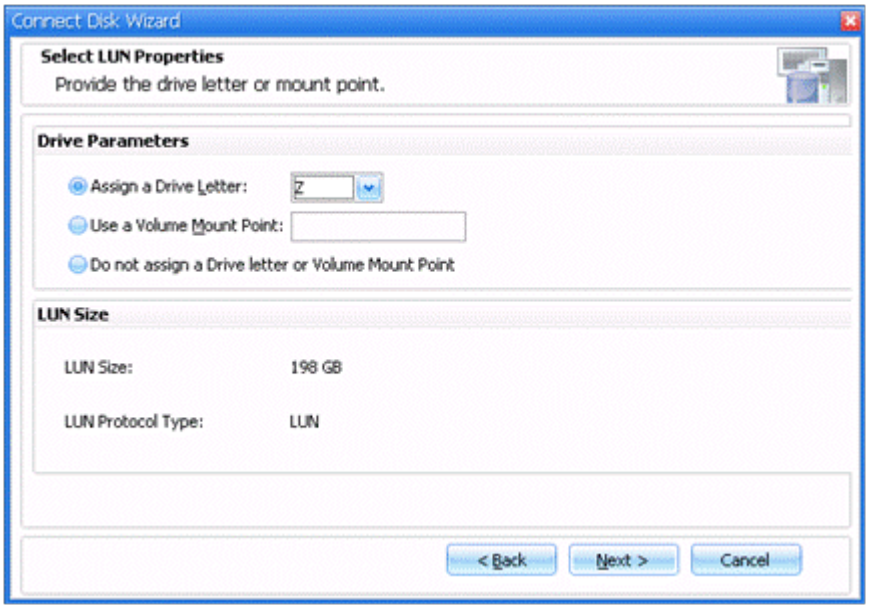

15. If you are using iSCSI, you are prompted to select an iSCSI initiator. Click the checkbox next to your initiator to select it.
16. Click Next.

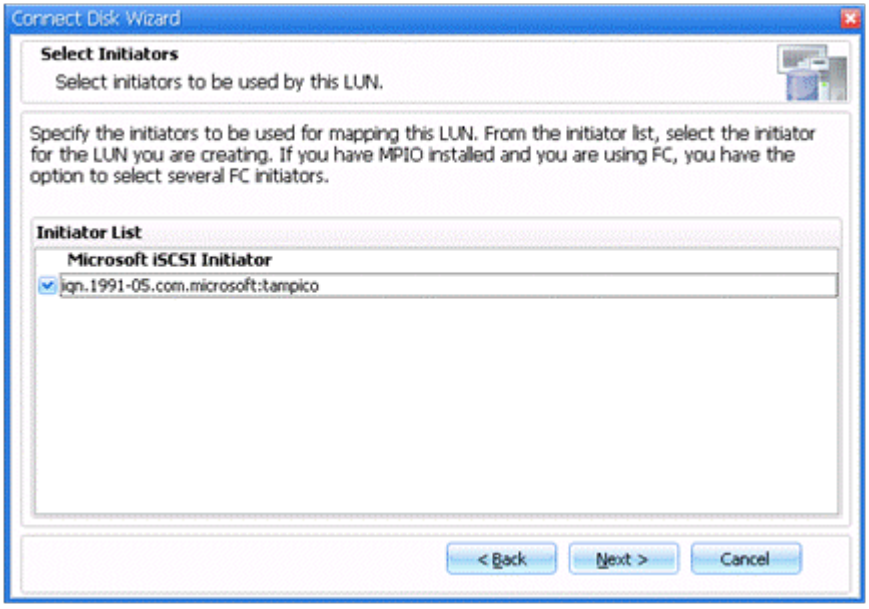

17. Select the Initiator Group management for iSCSI. The default selection is Automatic.
18. Click Next.

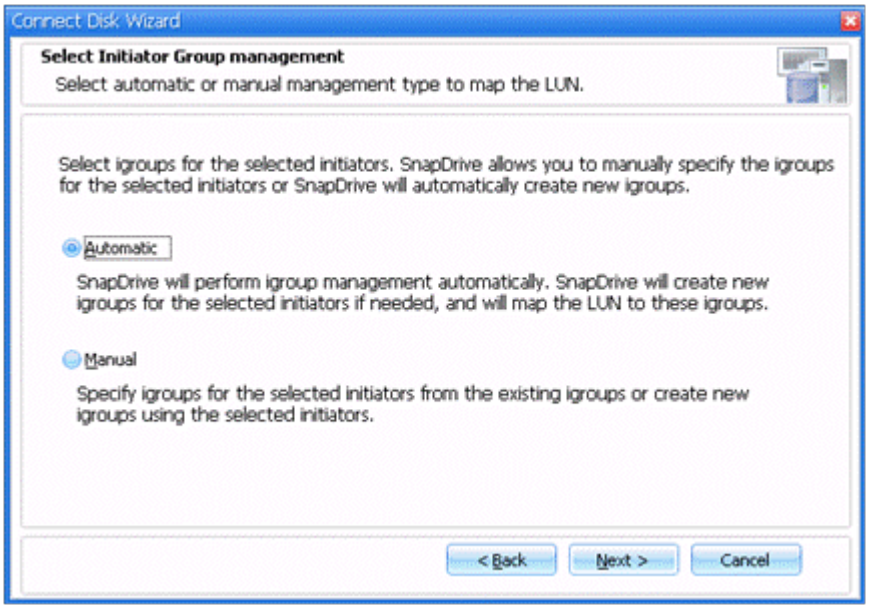

19. Review your settings on the Complete the Connect Disk Wizard page.
20. Click Finish.

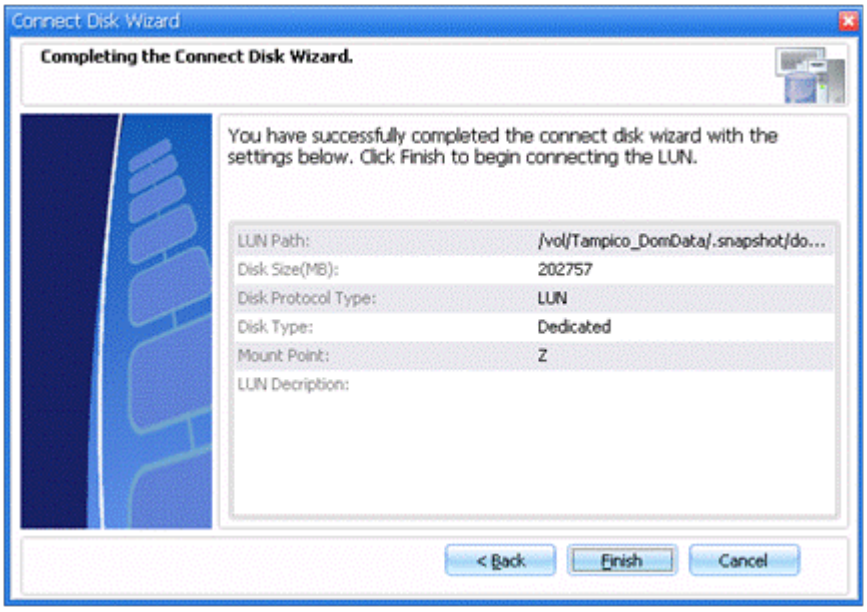

The Details section of the Disks page provides an update on the connection process.

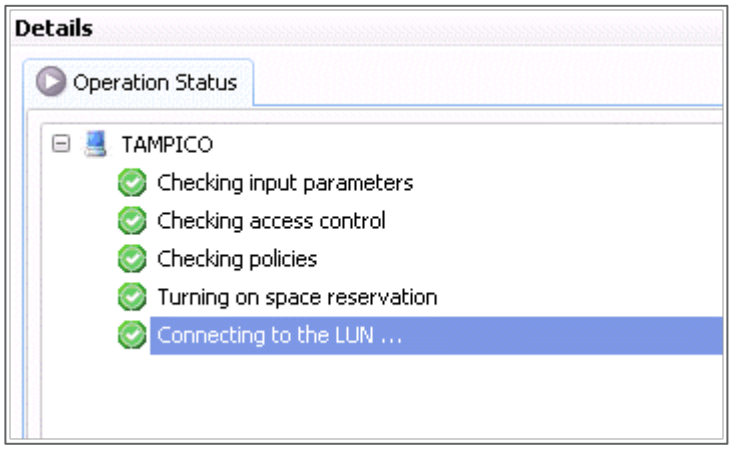

When the process is complete, the Details section changes from Operation Status to Disk Details and the details for the mounted disk are displayed.

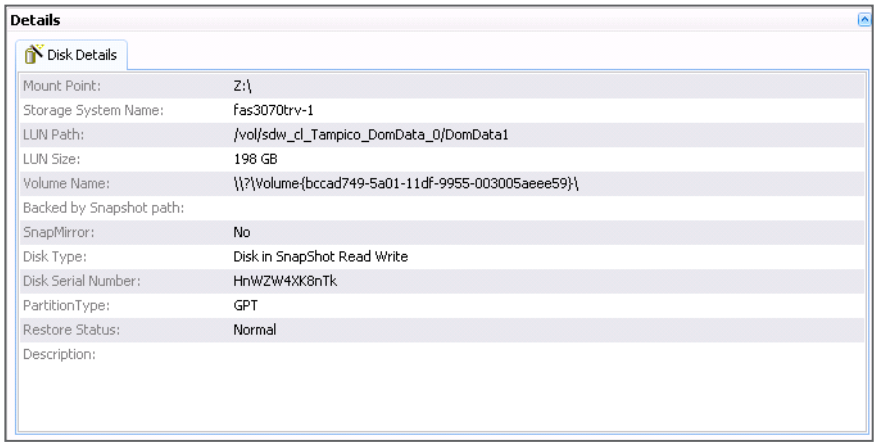

21. Leave the Computer Management window open so you can disconnect the disk later.
22. Open Windows Explorer® to browse to your Snapshot file system (Z:\).
23. Locate the database to restore and copy it to your workstation or to the active file system.

Note: To restore to an active file system, you will probably need to flush the database cache first by using the Domino CLI command `dbcache flush`.

A Domino administrator can open the restored database and copy a single e-mail from it into a user's production database.

Note: After copying the database to restore, remember to disconnect your Snapshot drive.

To disconnect your Snapshot drive:

24. Highlight the drive in Computer Management.
25. Right-click and select Disconnect Disk.

7 SNAPMANAGER FOR DOMINO AND SNAPVAULT

SnapVault can be used as an archiving method with SnapManager for Domino by using a script. SnapManager for Domino is typically used as a GUI for managing Domino backup and restores, but this method uses CLI commands to take SnapManager initiated Snapshot copies and to perform a SnapVault transfer.

Unlike some other SnapManager products, SnapManager for Domino does not have a Run Command After Operation option. Therefore, a script is necessary to process the SnapVault commands after SnapManager for Domino completes the backup operation. By using SnapManager for Domino in a script, the Domino APIs are still used to quiesce the databases for consistency.

7.1 ABOUT SNAPMANAGER FOR DOMINO

SnapManager for Domino is an integrated data management solution for Lotus Domino. SnapManager for Domino uses the Domino APIs to provide database consistency.

SNAPMANAGER FOR DOMINO REQUIREMENTS

Using SnapManager for Domino requires:

- Windows 2003/2008
- SnapDrive 4.x/6.x
- Domino 6.5.x–8.x
- Transactional logging enabled in the archived style
- Transactional logs located on a different volume than the volume for the Domino data

Note: Contact your NetApp sales representative or channel partner to find out how to purchase SnapManager for Domino or to get a trial version.

7.2 CONFIGURING SNAPMANAGER FOR DOMINO AND SNAPVAULT

Follow the directions up to section 5.2 to configure SnapManager for Domino and SnapVault. From that point forward, the steps are slightly different because the schedules are set by using Windows Task scheduler.

Schedule your retention policy by using the `@-` switch. When running the SnapVault `snap sched` command, the `@-` switch indicates a retention policy rather than a Snapshot schedule. Using the `@-`

switch indicates that SnapVault updates are driven by SnapManager rather than by the SnapVault Snapshot update schedule.

Execute the following command to create a SnapVault retention policy:

```
SnapVault snap sched [vol_name] [snapvault_name] [#snapshots to keep] @-
```

For example, to keep 30 Snapshot copies for the most recent SnapManager for Domino Snapshot volume, execute the following command:

```
SnapVault snap sched Athena_domdata dominosnap_Athena_domino_recent 30@-
```

Note: When using the retention policy `[@-],` the Snapshot names are slightly different on the SnapVault side. SnapManager for Domino retains the `recent` Snapshot copy, which is updated by using SnapVault each time a new Snapshot copy is taken with SnapManager for Domino. On the SnapVault end, a series of Snapshot copies are named with a suffix.

For example:

```
dominosnap_Athena_domino_recent.0 <newest>
dominosnap_Athena_domino_recent.1
dominosnap_Athena_domino_recent.2
dominosnap_Athena_domino_recent.3
dominosnap_Athena_domino_recent.4…
```

7.3 CREATE A SNAPMANGER FOR DOMINO AND SNAPVAULT SCRIPT

As part of a script to work with SnapVault, run SnapManager for Domino to make sure that the SnapVault updates are processed properly. This very simple script contains only three sections. This script runs from the host system; therefore, communicate with your storage system by using the remote shell (RSH) or the secure shell (SSH).

To coordinate SnapManager for Domino backups with SnapVault:

1. Execute the SnapManager for Domino backup command:

```
C:\Lotus\Domino\smd.exe –backup= [Domino partition] –del= [# of snapshots to
keep]
```

2. Update SnapVault:

```
ssh [secondary_filer] SnapVault update –s [snapshot]
[secondary_filer]:/vol/[vol]/[qtree]
```

3. Take a SnapVault Snapshot copy:

```
ssh [secondary_filer] SnapVault snap create [volume] [snapshot]
```

7.4 SAMPLE SCRIPT: SNAPMANAGER FOR DOMINO AND SNAPVAULT

The following sample script can be used in a SnapManager for Domino environment.

```
------------------------------------------------
REM SMD SnapVault Backup Script
REM Backup Domino keeping 10 snapshots
C:\Lotus\Domino\smd.exe --backup=1 --del=10
REM Update SnapVault secondary
```

```
rsh Secondary_Node SnapVault update -s dominosnap__Athena__Athena__recent
/vol/Athena_domdata/Athena_domdata
REM Create snapshot to archive the above transfers
rsh Secondary_Node SnapVault snap create Athena_domdata
dominosnap__Athena__Athena__recent
REM End sample script
---------------------------------------
```

8 CONCLUSION

Use SnapVault in your Domino environment to manage challenges relating to data protection and archiving for your Domino infrastructure. By using SnapVault, you can schedule frequent and efficient backups of your Domino environment, while minimizing media consumption and system overhead.

Compared to traditional backup methods, SnapVault offers a high-performance, cost-effective backup and recovery solution for NetApp storage systems. SnapVault enables you to accelerate backups, reduce media costs, and take the risk out of restores.

9 ACKNOWLEDGEMENTS

Any paper such as this is never solely the work of the author; it is the work of the collective team that helped the paper come together. I would like to acknowledge the efforts of Jeremy Merrill and Chris Blackwood, who have helped tremendously during the writing of this technical report. I would also like to thank Jawahar Lal and Michelle Nguyen for taking the time to review and revise this report.

NetApp provides no representations or warranties regarding the accuracy, reliability or serviceability of any information or recommendations provided in this publication, or with respect to any results that may be obtained by the use of the information or observance of any recommendations provided herein. The information in this document is distributed AS IS, and the use of this information or the implementation of any recommendations or techniques herein is a customer's responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. This document and the information contained herein may be used solely in connection with the NetApp products discussed in this document.

www.netapp.com