



Microsoft® SQL Server 2000: Providing Disaster Recovery with NetApp® SnapMirror® Technology

Gerson Finlev & Maryam Basiji | Network Appliance | TR 3108

TECHNICAL REPORT

Network Appliance, a pioneer and industry leader in data storage technology, helps organizations understand and meet complex technical challenges with advanced storage solutions and global data management strategies.

Table of Contents

1. Purpose and Scope
2. Description of NetApp SnapMirror Technology
3. General Assumptions and Requirements
4. Prerequisites
 - 4.1. Filer Prerequisites
 - 4.2. SQL Server Prerequisites
 - 4.3. Network Prerequisites
5. Setting Up a SnapMirror Process for SQL Server 2000 Data File Storage
 - 5.1. Step One
 - 5.2. Step Two
 - 5.3. Step Three
6. SnapMirror Software Operation
7. Disaster Recovery—Breaking the Mirror
8. Resyncing the Mirror
9. Conclusions
10. Support
11. Caveats
12. Revision History

[TR3108]

1. Purpose and Scope

This document describes techniques for setting up a disaster recovery configuration for Microsoft SQL Server databases using Network Appliance™ SnapMirror technology. Specifically, we cover the following issues:

- Description of the NetApp SnapMirror technology and its approach to disaster recovery
- The infrastructure required to support SnapMirror technology
- How to set up a SnapMirror for use in a Microsoft SQL Server environment
- How disaster recovery works in practice
- Resyncing the mirror following recovery from a disaster

Hardware and software used by the test environment are:

- Two NetApp F760s, Data ONTAP™ 5.3.7
- SQL Server 2000, Enterprise Edition
- Windows® 2000 Advanced Server

The methodology described in this paper shows how to recover a database to the same consistent state it was at when the SnapMirror shot was taken; committed transactions after the SnapMirror shot was taken are lost. If the requirement is to recover to the last committed transaction prior to the disaster, the database log file has to be located outside the filer on a medium that is accessible from the recovering system.

2. Description of NetApp SnapMirror Technology

NetApp SnapMirror technology provides asynchronous mirroring of data between filer volumes. Data on the source volume is periodically replicated to the target at a user-definable time interval,

with the range being from one minute to one month. At the end of each replication event, the mirror target volume becomes an exact block-for-block copy of the mirror source volume. At that point, the two volumes share identical data content and characteristics. The mirror is initialized by effectively copying the entire source volume to the target volume. Once this initial copy is complete, replication events thereafter copy only changed blocks from the source volume to the target volume. This provides a highly efficient data replication mechanism.

Architecturally SnapMirror software is a logical extension of the NetApp WAFL™ file system, particularly the Snapshot™ feature. Using Snapshots, you can create a read-only copy of an entire filer volume. This copy is made by essentially saving only changed blocks after a particular point in time. Two sequential Snapshots can then be compared and the differences identified. Since this comparison takes place at the block level, only the changed blocks need be sent to the mirror target. By implementing the update transfers asynchronously, data latency issues inherent with remote synchronous mirroring techniques are eliminated. The elegance of these two design features becomes particularly apparent when running mirror pairs over WAN topologies.

SnapMirror scheduling is configurable by the filer administrator, and since SnapMirror creates and uses its own Snapshots, is it not possible for the database server to know when a SnapMirror file will be updated. Therefore, if the requirement is to use the database in the mirror as a reporting database or to recover from soft database error (database objects deleted or modified by operational mistake) than another snapshot has to be used. Snapshots created on the source filer are automatically migrated to mirror filer, and SQL Server can attach to the database in the mirrored snapshot.

For more information on how Snapshots work, see [File System Design for an NFS File Server Appliance](#) by Dave Hitz, James Lau, and Michael Malcolm. Those with [NetApp On the Web](#), or "NOW," password access can also check the following documentation, which covers [Snapshots](#) in detail. ([NetApp On the Web](#) is the NetApp online support site. This site is open to all current Network Appliance customers.)

3. General Assumptions and Requirements

We assume that you are familiar with Microsoft SQL Server, the operation of NetApp filers, and the Windows operating system.

The examples in this technical report assume the following:

- The name of the source filer is **bg**.
- The name of the target filer is **lg**.
- The name of the administrative user account within SQL Server is **internal** and the password of this user is **microsoft**.
- All data is stored in the **mssql** directory on **vol2** on **bg**.
- All data is mirrored to **vol2** on **lg**.

4. Prerequisites

Following are the prerequisites for using SnapMirror on a NetApp filer with Microsoft SQL Server:

4.1. Filer Prerequisites

- Both target and source filers must be running the same Data ONTAP version. The respective software version must be 5.3 or higher.

- Both the SnapMirror and CIFS licenses must be enabled on all participating filers.
- An offline volume must be created or an existing volume must be taken offline. This volume will be the mirror volume. If using an existing volume, all premirror data on that volume will be overwritten once the baseline transfer has taken place. Because the contents of this volume will be overwritten by the SnapMirror volume data, this volume should not be the root volume.
- The source volume must not be a mirror and it must be online.
- The capacity of the mirror volume must be greater than or equal to the capacity of the source volume. The configuration of these volumes can differ geometrically (e.g., 16x18GB drives vs. 8x36GB drives), however, any geometric mismatch will result in a significant performance penalty.
- The source volume must be large enough to house the database(s) being mirrored.

4.2. SQL Server Prerequisites

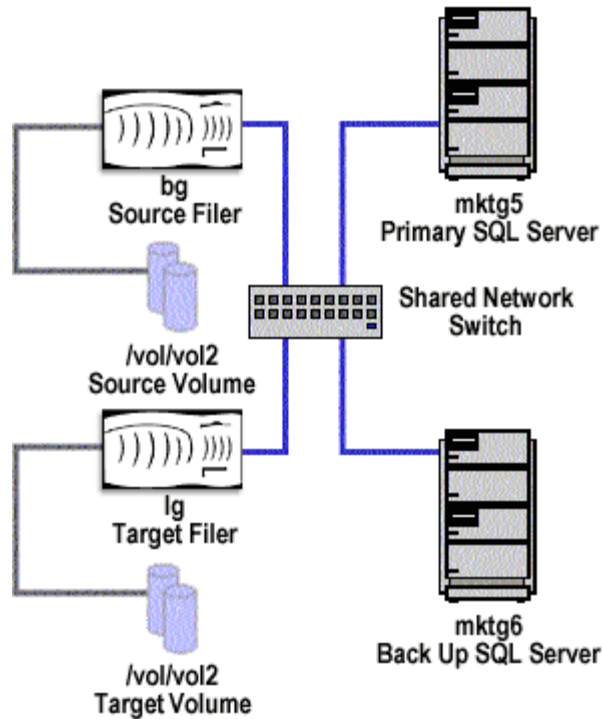
- A Microsoft SQL Server install on Windows Server is required to implement the exercises in this document.
- Both the SQL Server database data and log files must reside on the source volume.

4.3. Network Prerequisites

- A network connection is required between the SQL Server machine and the filer. 100BaseT, FDDI, and Gigabit Ethernet are all viable options. The faster the network, the better the performance.
- A robust network connection between the source filer and the target filer is required. It must have sufficient bandwidth to accommodate the anticipated data change rate and SnapMirror software overhead.
- The network connection type should be based on the following parameters:
 - Data transmission costs between the source and target filers
 - The source volume size
 - The data change rate
 - The SnapMirror software update schedule

In addition to the data change block transfers, each replication event requires an updated WAFL block map defined as 0.1% of the source volume size. The block map file is transferred for each mirror iteration regardless of the number of changed blocks. For example, a 250GB volume mirror will transfer a 250MB block map file plus all changed data blocks since the last Snapshot.

Following is the network connection we used to test this solution:



5. Setting Up a SnapMirror Process for SQL Server 2000 Data File Storage

Mirrors are as simple to set up and manage as filer volumes. Once all prerequisites have been met, the following must be done to initiate the process:

5.1. Step One (Implemented on the Source Filer):

Add the host name of the destination filer to the `/etc/snapmirror.allow` file on the source filer. In our test environment, `lg` is the name of the destination filer, so the `/etc/snapmirror.allow` file on the source filer (`bg`) appears as follows (this file can be viewed by mapping a drive to the root volume of the filer and looking in the `/etc` directory):

```
lg
```

Note: The filer is not shipped with a default `/etc/snapmirror.allow` file. Rather, it must be created using a text editor.

5.2. Step Two (Implemented on the Destination Filer):

Specify the following in the `/etc/snapmirror.conf` file on the destination filer:

- The source filer and volume
- The target filer and volume
- The maximum network bandwidth usage throttle
- A list of incremental update times in minutes, hours, days, and months

In our testing environment, the `/etc/snapmirror.conf` file on the destination filer `lg` appears as follows (see **Step One** for how to view this file):

```
bg:vol2 lg:vol2 kbs=5000 * * * *
```

The above text sets up the following SnapMirror software parameters:

- The **vol2** volume on **bg** will be replicated onto the **vol2** volume on **lg**
- A 5,000KB per second maximum throttle is set on the mirror
- The SnapMirror interval is set to occur every minute of every hour of every day of every week of every month

Note: The filer is not shipped with a default `/etc/snapmirror.conf` file. Rather, it must be created using a text editor.

5.3. Step Three (Implemented on Both the Source and Destination Filers):

At the command line of both the source and destination filers, enter the following command:

```
vol snapmirror on
```

It is important to note that the **vol snapmirror on** command does not persist across filer reboots. Rather, the command must also be put in the `/etc/rc` files of both source and destination filers. (Note: this command can be placed anywhere after the network interfaces are defined.) If placing the **vol snapmirror on** command in the `/etc/rc` file, following is how this file may appear (see **Step One** for how to view this file):

```
#Regenerated by registry Wed Apr 21 10:10:28 PDT 1999
#Auto-generated by setup Mon Apr 19 18:22:22 GMT 1999
hostname bg
ifconfig e0 `hostname` -e0 mediatype auto netmask 255.255.252.0
route add default 10.153.4.1 1
routed on
options dns.domainname 2700-1.netapp.com
options dns.enable on
...
vol snapmirror on
```

With SnapMirror turned on, the destination filer will read the `/etc/snapmirror.conf` file and, at the next scheduled mirror update, establish a connection with the source filer. If a baseline version of the mirror does not exist, the filer takes a Snapshot of the source volume and transfers all the data in the Snapshot from the source volume to the mirror. If a **vol status** command is issued at this point, the following result will appear (noting that, in this example, **vol2** is our mirror volume):

Command:

```
bg>vol status
```

Result:

```
Volume State Status Options
vol0  online normal  root, nosnap=on
vol1  online normal
vol2  online snapmirrored
```

Subsequent updates to the mirror are made according to the schedule specified in the `/etc/snapmirror.conf` file.

6. SnapMirror Software Operation

SnapMirror has two distinct phases: initialization and incremental update. The initialization phase consists of a level 0 replication event in which a Snapshot is created on the source volume and then sent in its entirety to the target volume. The amount of time required to replicate the entire source volume over to the target volume depends on many factors, including the network connection. For example, a 250GB mirror initialization that would take approximately 7.5 hours to

complete across a 100BaseT full-duplex link may take 1.5 hours across a gigabit link. The level 0 event serves to initialize or "seed" the mirror volume since it contains every block in the source volume as of the time the Snapshot was created.

After mirror initialization is complete, the target filer(s) examines its **/etc/snapmirror.conf** file every minute to see if there are any scheduled updates. This allows for modification of the mirror's configuration without disrupting the mirror. When an incremental update schedule time is due, a new Snapshot is taken and compared to the previous Snapshot. The different blocks and block map file are sent to the mirror target. In contrast to the level 0 initialization, the data mirrored is typically much smaller. Note that at all times the mirror target file system is in a consistent state.

Several special cases should be noted:

- If an initialization (level 0) replication event is interrupted for more than nine minutes (e.g., a network outage), it will abort. Partial level 0 events are not recoverable, so the process must be restarted.
- The KB per second maximum throttle parameter in **/etc/snapmirror.conf** can be changed at any time, and the edited values will take effect within two minutes. The exception to this is mirror initialization where a bandwidth throttle is already in effect. In this case, the bandwidth throttle cannot be modified until the initialization phase is complete or the process is interrupted.
- Incremental updates will not start until the level 0 initialization is complete.
- Any incremental updates that are missed are simply skipped. Subsequent incremental updates transmit any skipped data, so no data is lost.
- Incremental updates in progress will run to completion according to the configuration settings and available network bandwidth. If a new incremental update is scheduled to start while an existing update is in progress, it is considered a schedule miss and this is skipped.
- The SnapMirror process can be stopped and restarted at any time by issuing the following command sequence on either filer:

```
vol snapmirror off  
[arbitrary time interval]  
vol snapmirror on
```

As long as the target volume remains read-only during the time between turning the SnapMirror process off and on, the mirror remains intact.

- To break the mirror, SnapMirror must be turned off, at which time the target volume will be placed into read/write mode. In order to turn SnapMirror off, the following command should be executed on the destination filer (in our testing environment, **vol_name** is **vol2**):

```
vol options vol_name snapmirrored off
```

Note: The SnapMirror file can only be broken if a SnapMirror event is not in process.

7. Disaster Recovery—Breaking the Mirror

After creating a SnapMirror relationship between two filers we implemented the following test:

1. We created the following table on the database **test_db**. The database's data and log files were located in the **mssql** directory on **vol2** on **bg**:

```
DROP TABLE hammer_tab;

CREATE TABLE hammer_tab
  ( value INT NOT NULL );
```

2. We then created the following stored procedure in the table **hammer_tab**:

```
USE test_db
GO;
CREATE PROCEDURE add_data
AS
DECLARE @counter INT
SELECT @counter = 1
WHILE ( @counter < 100000000000 )
BEGIN
  INSERT INTO hammer_tab
    ( value )
  VALUES
    ( @counter );
  PRINT 'Value is: ' + convert ( varchar(6), @counter1 ) SELECT @counter = @counter + 1
END
```

3. We then executed **add_data** from the SQL Server's Query Analyzer. With the procedure running, we allowed several SnapMirror events to occur.
4. Toward the end of the ten-minute procedure run, we changed the SnapMirror interval in the **/etc/snapmirror.conf** file to a less frequent mirror. We did this because we were about ready to simulate a filer/SQL Server failure followed by a recovery of both servers. We did not want a SnapMirror event to initiate once the filer recovery had completed. If an event was triggered, we would not be able to break the mirror.
5. Next we simulated the above-mentioned disaster with a goal of mimicking a complete loss of service of both the filer and the database server. We did this in the following manner:
 - o Failed the source filer (**bg**) by powering it down
 - o Failed the SQL Server by turning off the MSSQLServer service while the query was still running

The failure of both the filer and the SQL Server resulted in the following activity from within the Query Analyzer (note that the last row inserted has a key value of 150942):

```
(1 row(s) affected)
Value is: 1
(1 row(s) affected)
Value is: 2
(1 row(s) affected)
Value is: 3
(1 row(s) affected)
Value is: 4
... Many more like this ...
(1 row(s) affected)
Value is: 150939
(1 row(s) affected)
```



```
Value is: 150940
(1 row(s) affected)
Value is: 150941
(1 row(s) affected)
Value is: 150942
```

```
[Microsoft] [ODBC SQL Server Driver] [Named Pipes] ConnectionRead (GetOverLappedResult()).
[Microsoft] [ODBC SQL Server Driver] [Named Pipes] Connection broken.
```

Connection Broken

6. After changing the mirror interval, we broke the mirror on **vol2** (making lg's **vol2** read/write) by issuing the following command on **lg**:

```
vol options vol2 snapmirrored off
```

7. We then created a CIFS share (**sqlldb**) pointing to the directory (**mssql**) on the SnapMirror volume (**vol2** on **bg**) in which the data resides.
8. After creating the share, we created a new SQL Server database (**test_db_new**) that pointed to the data and log files that were on **lg**'s now read/write volume **vol2**. We did this by executing the following statement from an ISQL prompt:

```
dbcc traceon ( 1807 )
go
sp_attach_db 'test_db_new', '\\lg\sqlldb\test_db_data.mdf', '\\bg\sqlldb\test_db_log.ldf'
go
dbcc traceoff ( 1807 )
```

9. We then queried the **hammer_tab** table in both **test_db** and **test_db_new** to find out if the data in the table on each database matched. Following are the query scripts and their respective result sets:

Query on test_db:

```
USE test_db
SELECT * FROM hammer_tab
```

Result set from query on test_db:

```
value time
1
2
3
4
... Many more rows incremented by one ...

150939
150940
150941
150942
(150942 row(s) affected)
```

Query on test_db_new:

```
USE test_db_new
SELECT * FROM hammer_tab
```

Result set from query on test_db_new:

```
value time
1
2
3
4
... Many more rows incremented by one ...
130941
130942
130943
130944
(130944 row(s) affected)
```

Note that the key value for the source database matches precisely to that which was reported to the client. This is consistently our experience: When a NetApp filer experiences a dirty shutdown, no transactions are lost, although the SQL Server may apply a recovery to the database on the source filer (**bg**). However, the target database did lose approximately 20,000 rows over the ten minutes that the test was run. This is to be expected. The mirror interval was set to 60 seconds, and the mirror relies on consistency points. Thus, the mirror can be as much as two minutes out-of-date with the interval set in this way. (A longer interval would result in more lost transactions, but also lower overhead.) In our example, a total 150,942 rows were processed on the source database table. This is a rate of 15,094.2 rows per minute. Given that the mirror can be out-of-date by minutes, the number of rows that did not mirror over to the target database table could have been as high as 30,188.4 (20% of the table). The bottom line is that the target database can be created as a mirror of the source database, with a loss of only a few very recent transactions. For a disaster recovery solution, this is perfectly acceptable to the vast majority of customers.

8. Resyncing the Mirror

Currently the ability to restart a broken mirror is not available, although this feature is planned for a future release. Presently, the method for reestablishing a mirror is to reinitialize the mirror. If the target volume has been brought online and then written to, this may need to be done twice.

Once to replicate the data back to the source volume. (In this case the mirror relationship between the two filers is flipped.)
Once to reinitialize with the mirror relationship back to normal.

An alternative is to back up the target volume's data to tape, and then restore that data onto the source filer. This avoids the need to reinitialize the mirror twice; however, you currently cannot "seed" a mirror from tape. This is also planned for a future release.

9. Conclusions

The NetApp SnapMirror technology provides compelling advantages for the SQL Server DBA seeking to support disaster recovery of a mission-critical SQL Server database. Specifically:

- The network traffic generated by the mirror process can be throttled to allow support for WAN connections.
- The mirror interval is user-configurable and can be changed on-the-fly.
- In the event of a disaster on the source side of the mirror, the target side can be easily brought online, and the SQL Server database can return to normal use after a brief recovery process.

- The near real-time nature of SnapMirror software means that the source filer will still operate normally when the network link to the target filer is broken. This is not true of synchronous mirroring, where the mirror link is essential to continued operation of the source machine.

Using SnapMirror technology, the DBA can assure the customer that the only transactions that will be lost are those that occur between the mirror events.

10. Support

As a result of discussion between Microsoft and Network Appliance, Microsoft created dbcc trace flag 1807 to enable the storage of SQL Server databases on Network Appliance filers. This flag is required to successfully migrate or create SQL Server 7.0 or SQL Server 2000 databases on a filer. Though databases can be placed on any network-attached device using this flag, Microsoft specifically designed the flag for Network Appliance. Network Appliance fully supports the specific configurations described in this document. If you experience any problems implementing the techniques described in this technical report, please contact Network Appliance's [SQL Server Technical Marketing Engineer](#).

11. Caveats

Though this document provides a comprehensive overview of SnapMirror technology, additional reading may be beneficial. The Data ONTAP **Systems Administrator's Guide** that accompanies your software version is an excellent resource.

12. Revision History

Date	Name	Description
3/22/2001	Gerson Finlev	Update TR3076 to SQL Server 2000 functionality
6/12/2001	Maryam Basiji	Update



Network Appliance, Inc.
495 East Java Drive
Sunnyvale, CA 94089
www.netapp.com

© 2005 Network Appliance, Inc. All rights reserved. Specifications subject to change without notice. NetApp, NetCache, and the Network Appliance logo are registered trademarks and Network Appliance, DataFabric, and The evolution of storage are trademarks of Network Appliance, Inc., in the U.S. and other countries. Oracle is a registered trademark of Oracle Corporation. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.