

HP OpenView

Storage Mirroring application notes

High availability for Windows SharePoint Services

Legal and notice information

© Copyright 2005-2006 Hewlett-Packard Development Company, L.P.

Hewlett-Packard Company makes no warranty of any kind with regard to this material, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. Hewlett-Packard shall not be liable for errors contained herein or for incidental or consequential damages in connection with the furnishing, performance, or use of this material.

This document contains proprietary information, which is protected by copyright. No part of this document may be photocopied, reproduced, or translated into another language without the prior written consent of Hewlett-Packard. The information is provided "as is" without warranty of any kind and is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Microsoft, Windows, Windows NT, and Windows XP are U.S. registered trademarks of Microsoft Corporation.

Storage Mirroring High availability for Windows SharePoint Services application notes

Document overview

This document is a Storage Mirroring application note. An application note provides guidelines on the use of Storage Mirroring in a specific environment.

This document contains:

- **Document overview**—Explains what an application note contains, how it should be used, what you need to know before trying to use the application note, and where you can go for more information.
- **Solution overview**—Explains how the application works with Storage Mirroring and describes the considerations that you must weigh when implementing your Storage Mirroring solution. Review this section to make sure that you understand the theory involved with using Storage Mirroring and your application. Includes both basics, such as system requirements, as well as configuration and environment-specific topics, such as interactions with specific clients or special considerations for WAN (Wide Area Network) environments. Pay special attention to those topics that are directly related to your environment.
- **Sample implementations**—Describes specific examples of how to use Storage Mirroring for this solution. This includes information about the specific system setup used in the sample implementation. Use these procedures as a guideline for creating your own implementation. Because no two environments or configurations are exactly the same, you will probably need to implement additional or different steps than what is documented here in order to make the solution work in your environment.

Audience

This document is written for network and application administrators who have a working understanding of the applications and environments where the Storage Mirroring solution is to be deployed. You may need to expand on the documented information in order to customize the solution to fit your environment.

Before you use this application note, you should have an understanding of:

- Storage Mirroring
- SharePoint Services

Expectations

Application notes are intended to provide a framework for configuring a Storage Mirroring solution in a specific environment and to draw attention to decisions you will need to make when configuring your solution.

Because there are an infinite number of possible configuration, network, and environment scenarios, application notes contain general configuration guidelines as well as an example configuration procedure that has been tested for a specific environment.

This document assumes that you are comfortable working with your operating system, Storage Mirroring, and SQL Server.

Related documentation

Before you begin to configure your solution, make sure that you have complete documentation for your operating system, application, and Storage Mirroring. This application note does not provide step-by-step instructions for using standard operating system, application, and Storage Mirroring functionality.

The following documents contain additional information that you may need while setting up this solution:

- *HP OpenView Storage Mirroring user guide* or online documentation
- Reference guides or documentation for SharePoint Services

Getting help

Hewlett-Packard has application notes that describe how to configure Storage Mirroring with a variety of popular third-party applications. These application notes are available on the Storage Mirroring web site: <http://h18006.www1.hp.com/products/storage/software/sm/index.html>.


For help using Storage Mirroring, refer to the Storage Mirroring online manual or online help.

Solution overview

Windows SharePoint Services is a collection of Windows Server 2003 services that allows you to create team-based web sites for document and information sharing, management, and coordination. Storage Mirroring provides real-time enterprise data protection and replication. Storage Mirroring can be used to provide high availability for Windows SharePoint Services.

This document describes the steps necessary to configure Storage Mirroring to provide high availability for Windows SharePoint Services using Microsoft SQL Server 2000. These procedures allow a secondary server to assume the identity and role of a failed server while maintaining the availability of services with minimal disruption or data loss.

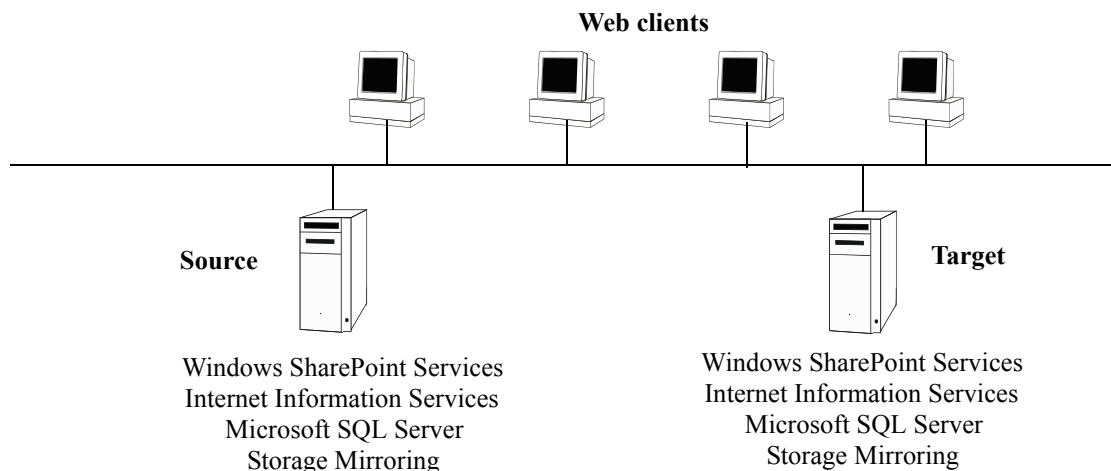
To complete these instructions, you will install Windows SharePoint Services, SQL Server 2000, and Storage Mirroring. You will configure Storage Mirroring for replication and failover.

 **NOTE:** Due to the complexities of these applications, this document is intended for network administrators with experience installing, configuring, and maintaining network applications including Storage Mirroring, Microsoft SQL Server, and Windows SharePoint Services.

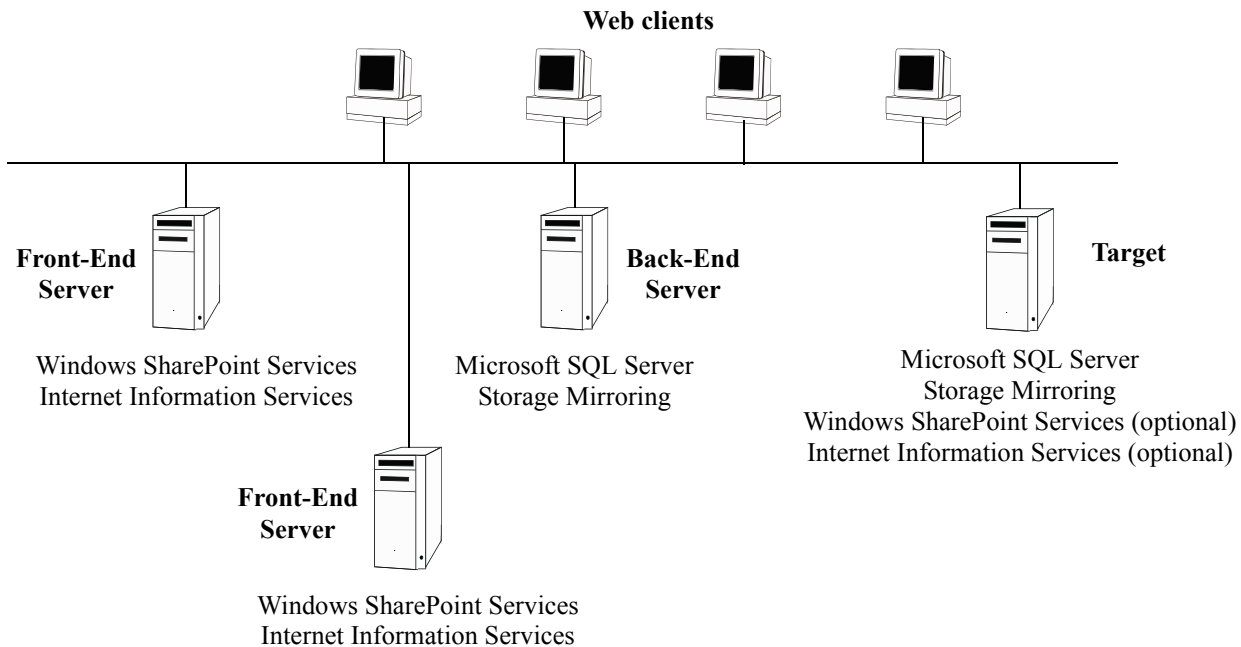
Configurations

Windows SharePoint Services (WSS) combines stateless front-end web servers and back-end database servers. It can be configured in two different ways.

- **Single Standalone Server**—In a standalone SharePoint environment, a single server is configured to run Windows SharePoint Services, Internet Information Services (IIS), and Microsoft SQL Server. Storage Mirroring can be used to replicate and failover the SQL server database to a Storage Mirroring target machine, and WSS and IIS can be pre-configured on that target machine to continue processing requests in the event the primary server fails.



- **Server Farm**—In a server farm configuration, one or more servers are configured as front-end web servers, running both WSS and IIS, accessing data from a SQL Server database located on an independent server. Again, Storage Mirroring can be used to replicate and failover the SQL server to an independent target server. Since a server farm configuration uses multiple front end WSS servers, usually in some type of load-balancing configuration with its own failover capability (such as Windows Load Balancing), it is usually only necessary to failover the SQL databases storing the data so that surviving front-end WSS servers can continue to access the database. However, just as in the single standalone server case, it is also possible to configure a WSS front-end server to run on the StorageMirroring target as well, if desired, ensuring that WSS can continue to operate even with a singleModifying the sample script files



After you modify the sample scripts, save them with a new name to remove the `SAMPLE_` prefix. Copy the scripts to the directory where Storage Mirroring is installed.

The sample batch files provided are only examples. Because no two environments or configurations are exactly the same, you **MUST** modify the sample scripts in order to make the solution work in your environment.

Modifying the sample script files

After you modify the sample scripts, save them with a new name to remove the `SAMPLE_` prefix. Copy the scripts to the directory where Storage Mirroring is installed.

The sample batch files provided are only examples. Because no two environments or configurations are exactly the same, you **MUST** modify the sample scripts in order to make the solution work in your environment.

Configuring the replication set

For a default SQL installation, you *must* include the SQL Data directory in your replication set (`<drive>:\Program Files\Microsoft SQL Server\MSSQL\Data`). This directory includes the SQL application data and transaction log files.

In addition, you *may* need to include the following directories, depending on your environment configuration:

- `Tempdb` files. You should always include the `tempdb` files, unless you can determine that they are not being used by any application. Some applications, including Prophecy, PeopleSoft, and BizTalk, write data to the `tempdb` file.
- Any other directories (even if on different drives) that you may have created to store SQL data files.

- The SQL Log directory (<drive>:\Program Files\Microsoft SQL Server\MSSQL\Log). This directory includes the SQL error logs from the source, which may be useful when determining the cause of a source failure. In order to prevent overwriting the target's error logs, you will need to replicate the log files to a different location on the target. Create a separate replication set for the error logs and select the "All to One" mapping option.

You *do not* need to replicate the application files, since they already exist on the target machine.

Enabling compression

By enabling compression, you can reduce the amount of bandwidth needed to transmit Storage Mirroring data. When compression is enabled, the data is compressed before it is transmitted from the source. When the target receives the compressed data, it uncompresses it and then writes it to disk. On a default Storage Mirroring connection, compression is disabled.

Because the files that should be included in a SQL replication set can generate a significant amount of data, you should enable compression for the connection. For more information about enabling compression, see the *HP OpenView Storage Mirroring user guide*. However, keep in mind that the process of compressing data impacts processor usage. If you notice an impact on performance while compression is enabled in your environment, either adjust to a lower level of compression, or leave compression disabled.

Configuring SQL memory usage

Storage Mirroring uses memory to queue operations and data on both the source and target. Since the source server is typically running a production application, it is important that the amount of memory Storage Mirroring and the other applications use does not exceed the amount of RAM in the system. If the applications require more memory than there is RAM, the system will begin to swap pages of memory to disk and the system performance will degrade.

For instance, SQL Server will use all of the available system memory when needed by default, and it may use almost all of the system memory during high-load operations. These high-load operations are precisely what cause Storage Mirroring to need memory to queue the data being changed by SQL Server. On a server with 1GB of RAM running SQL Server and Storage Mirroring, you might configure SQL Server to use only 512MB and Storage Mirroring to use 256MB, leaving 256MB for the operating system and other applications on the system. Many other server applications will use almost all system memory by default, so it is important to check and configure applications appropriately, particularly on high-capacity servers.

Sample Implementation

This section describes an example of how to figure Storage Mirroring and SharePoint Services. Use these procedures as a guideline for creating your own implementation. Because no two environments or configurations are exactly the same, you will probably need to implement additional or different steps than what is documented here in order to make the solution work in your environment.

Requirements

Each server must meet the following requirements:

- A licensed copy of Microsoft Windows Server 2003
- A licensed copy of Microsoft SQL Server 2000 with Service Pack 3A or later

NOTE: Although SharePoint Services will work either with Microsoft SQL Server 2000 Desktop Engine (MSDE 2000) or the full version of SQL Server, this solution only works with the full version of SQL Server.

If your SharePoint Services system is running on MSDE, the database must be upgraded to the full version of SQL Server (using the migration procedure recommended by Microsoft) before the system can be protected by Storage Mirroring.

-
- A licensed copy of Storage Mirroring with the latest service pack

Preparing the source(s)

NOTE: If you are using a server farm configuration with multiple front-end WSS servers, you must have a mechanism in place to redirect the clients to alternative front end front-end WSS servers in the event that the front-end WSS server (to which a client is connected) fails. Typically, this is done through some type of load-balancing configuration with its own failover capability (such as Windows Load Balancing).

1. Configure each source as a Windows member server.
2. Install Microsoft SQL Server on each source that will be used as a back-end server.
3. Install Windows SharePoint Services on each source that will be used as a front-end server.
4. Record the following WSS information for each front-end server.

Front-End Windows SharePoint Services Server	Application Pool Names and Settings	Account Mode	Configuration and Content Database Names	Site Administrator Name and Credentials

5. If you have any web sites created outside of WSS that will be integrated with WSS, record that information.

Additional Web Site Information

6. Record the following SQL information for each back-end server.

Back -End SQL Server	Windows SharePoint Services Configuration and Content Database Name	Windows SharePoint Services Configuration and Content Database Path

7. Install Storage Mirroring on each source server using the installation defaults. See the *HP OpenView Storage Mirroring getting started guide* for details.

Preparing the target(s)


1. Configure each Windows 2003 target as a member server in the same Active Directory domain or trusted domain environment as the source.
2. Install Storage Mirroring on each target server using the installation defaults. See the *HP OpenView Storage Mirroring getting started guide* for details.
3. The remaining target preparation steps are dependent on what role your target must take on in the event of a source failure.
 - If you want your target only to function as a SQL back-end server, complete the steps in “[SQL target](#)” on page 9.
 - If you want your target to function as both a WSS/IIS front-end server and a SQL back-end server, complete the steps in “[WSS/IIS and SQL target](#)” on page 9.

SQL target

1. Install SQL on the target using the same configuration as the source.
2. Disable the SQL server auto-start policies. Verify that SQL Server and SQL Server Agent are not set to auto-start.
3. Stop the following services and set them to manual startup so that all SQL files are closed on the target so that the source can replicate the changes. Make sure that you set all of the SQL services that are relevant to your environment.
 - Distributed Transaction Coordinator
 - MSSQLServer
 - SQLServerAgent
 - Microsoft Search

WSS/IIS and SQL target

1. Install Windows SharePoint Services on the target. After the installation, the installation wizard will automatically start the configuration process for you. You will be using the information you recorded earlier to configure the target identically to the source.

 **NOTE:** If you need to launch the configuration process manually, select **Start, Administrative Tools, SharePoint Central Administration**.

2. Under **Server Settings**, select **Set Configuration Database Server**.
3. For **Database server**, enter the name of the back-end SQL server.
4. For **SQL Server database name**, enter the source server configuration database name.
5. Enable **Connect to existing configuration database** and click **OK**.
6. Under **Virtual Server Configuration**, select **Extend or upgrade virtual server**.
7. Select your target from the **Virtual Server** list.
8. Under **Provisioning Options**, select **Extend and map to another virtual server**.
9. Set the **Server Mapping** and **Application Pool** settings to the same as those used on the source.
10. Click **OK**.
11. At this point, your configuration of the WSS/IIS portion is complete. If desired, you can use this copy of WSS/IIS to access the SQL database on the source. If you prefer to keep this copy in a dormant state, only to be used if the source front-end fails, you must set the following services on the target to manual startup.
 - IIS Admin
 - World Wide Web Publishing Service
 - Simple Mail Transfer Protocol (SMTP)
 - Network News Transfer Protocol (NNTP)
 - HTTP SSL (Secure Sockets Layer)

- SharePoint Timer Service
12. Install SQL on the target using the same configuration as the source.
 13. Disable the SQL server auto-start policies. Verify that SQL Server and SQL Server Agent are not set to auto-start.
 14. Stop the following services and set them to manual startup so that all SQL files are closed on the target so that the source can replicate the changes. Make sure that you set all of the SQL services that are relevant to your environment.
 - Distributed Transaction Coordinator
 - MSSQLServer
 - SQLServerAgent
 - Microsoft Search

Configure and begin mirroring and replication

The following steps create the Storage Mirroring replication set and establish the connection between the source and target. You will need to repeat these steps if you have multiple sources and/or targets.

1. On the source, select **Start, Programs, Storage Mirroring, Management Console**.
2. Double-click your source machine to log on.
3. Right-click the source and select **Properties**.
4. On the Source tab, enable **Block Checksum All Files on a Difference Mirror** and click **OK**.
5. Right-click the source machine and select **New, Replication Set** and enter the desired name for the replication set.
6. Select the SQL data you wish to protect. Most likely, this will only include the SQL data and log files. It is not necessary to replicate the application files since they already exist on the target machine.
 - Select the SQL data you wish to protect. For more information about what you should include in your SQL replication set, see "[Configure and begin mirroring and replication](#)" on page 10.)
7. Right-click the replication set name and select **Save** to save the replication set.
8. Drag and drop the replication set onto the target. The Connection Manager dialog box opens.
9. The **Source Server, Target Server, Replication Set**, and **Route** fields will automatically be populated. If you have multiple IP addresses on your target, verify the **Route** field is set to the correct network path. (For detailed information on connecting a source and target, see the *HP OpenView Storage Mirroring user's guide*.)
10. Select **One to One** to map the replication set data from the source to an identical volume/directory structure on the target.
11. On the Orphans tab, select the **Move/Delete Orphan Files** checkbox.
12. On the Mirroring tab, select the type of mirror, either **Full** or **File Differences**, to perform.

NOTE: If the target has been previously mirrored to or restored, select File differences with the Use block checksum option so that only the changed data is sent across the network.

13. Click **Connect** to start the mirror and replication processes.

NOTE: If you start Windows SharePoint Services and mount the replicated databases on the target, or if the data on the target is otherwise modified, the data on the source and target will no longer match. If the updated data on the target is not needed, perform a full or difference with block checksum mirror from the source to the target. If the updated data on the target is needed, restore the data from the target to the source.

Your data is protected after the mirror is complete and the Mirror Status has changed to **Idle**.

Configure failover and begin failure monitoring

The following steps should be performed on the target machine.

1. If a failure occurs, you will want to have the appropriate services start on the target automatically. To do this, create a batch file on the target called `postover.bat` using the sample batch file below. Save the batch file to the same directory on the target where your Storage Mirroring files are installed.

SAMPLE POSTOVER.BAT

```
rem Sample post-failover script for Windows SharePoint Services and SQL 2000.
rem You may need to remark out some commands depending on the function of your target.

rem The following commands should be used if your target is functioning as a back-end SQL target.
rem Make sure that you start all of the services, SQL included, relevant to your
rem environment.
net start "Distributed Transaction Coordinator"
net start "MSSQLServer"
net start "SQLServerAgent"
net start "Microsoft Search"
net start "MSSQLServerADHelper"

rem The following commands should be used if your target has WSS/IIS loaded, but it is in a dormant
rem state. In this case, the WSS/IIS services will start and the target will function as a front-end
rem WSS/IIS server.
net start "IISAdmin"
net start "World Wide Web Publishing Service"
net start "Simple Mail Transfer Protocol (SMTP)"
net start "Network News Transfer Protocol (NNTP)"
net start "HTTP SSL"
net start "SharePoint Timer Service"
```

2. After a failure is resolved, you will be ready to bring your source back online. At this time, you will want to stop the appropriate services on the target automatically. To do this, create a batch file on the target called `preback.bat` using the sample batch file below. Save the batch file to the same directory on the target where your Storage Mirroring files are installed.

SAMPLE PREBACK.BAT

```
rem Sample pre-failback script for Windows SharePoint Services and SQL 2000.
rem You may need to remark out some commands depending on the function of your target.
rem The following commands should be used if your target is functioning as a back-end SQL target.
rem Make sure that you start all of the services, SQL included, relevant to your
rem environment.
net stop "Distributed Transaction Coordinator"
net stop "MSSQLSERVER" /y
net stop "SQLSERVERAGENT"
net stop "Microsoft Search"

rem The following commands should be used if you want your target to revert back to a dormant WSS/IIS
rem state.
net stop "IISAdmin" /y
net stop "SharePoint Timer Service"
```

3. Select **Start, Programs, Storage Mirroring, Failover Control Center**.
4. Select the target machine from the list of available machines. If the target you need is not displayed, click **Add Target**, enter the machine name, click **OK**, and then click **Login**.
5. To add a monitor for the selected target, click **Add Monitor**. Type the name of the source machine and click **OK**. The Monitor Settings window will open.
6. In the Monitor Settings window, mark the IP address that is going to failover and verify that **Adding Source Identity to Target** is selected.
7. Under **Active Directory**, enable **Failover Hostname** and **Failback Hostname**.
8. Click **Account** and specify a username and password with full domain administrative privileges.
9. Highlight the machine name under Names to Monitor and click **Scripts**. Specify the scripts that were created earlier using `postover.bat` for the target post-failover script and `preback.bat` for the target pre-failback script.
10. Click **OK** to go back to the Monitor Settings dialog box and then click **OK** to begin monitoring the source machine.

In the event of a source machine failure, your target machine is now ready to stand in for your source. For information on monitoring failover, see the *HP OpenView Storage Mirroring user's guide*.

Restoring your data


If your source experiences a failure, such as a power, network, or disk failure, your target machine will stand in for the source while you resolve the source machine issues. During the source machine downtime, data is updated on the target machine. When your source machine is ready to come back online, the data is no longer current and must be updated with the new data on the target machine.

1. Verify that your source machine is not connected to the network. If it is, disconnect it.
2. Resolve the source machine problem that caused the failure.


 **NOTE:** If you must rebuild your hard drive, continue with step 3.

If you do not need to rebuild your hard drive, verify that the Storage Mirroring connection on the source has been disconnected (right-click the connection in the Storage Mirroring Management Console and select **Disconnect**) and then continue with step 6.

-
3. Install Windows. Since your source machine is not connected to the network, go ahead and use the source's original name and IP address.
 4. Install Storage Mirroring using the installation defaults.
 5. Install SQL using the configurations that you recorded when originally configuring the source.
 6. On the source, stop the following services so that all SQL files are closed on the source and the Storage Mirroring target can restore the changes back to the source. Make sure that you stop all of the services, SQL included, relevant to your environment:
 - Distributed Transaction Coordinator
 - MSSQLServer
 - SQLServerAgent
 - Microsoft Search
 - MSSQLServerADHelper
 7. On the target, select **Start, Programs, Storage Mirroring, Failover Control Center**.
 8. Select the target machine that is currently standing in for the failed source.
 9. Select the failed source and click **Failback**. The pre-failback script entered during the failover configuration stops the SQL services on the target so that no additional changes can be made.
 10. You will be prompted to determine if you want to continue monitoring the source server. Do not choose **Continue** or **Stop** at this time.
 11. Connect the source machine to the network.
 12. After the source is back online, select whether or not you want to continue monitoring this source machine (**Continue** or **Stop**).

 **NOTE:** Verify that the Storage Mirroring connection on the source has been disconnected (right-click the connection in the Storage Mirroring Management Console and select **Disconnect**).

-
13. To begin the restoration process, open the Storage Mirroring Management Console and select **Tools, Restoration Manager**.

 **NOTE:** You can also run the Storage Mirroring DTCL automated restoration script, which can be found in the *HP OpenView Storage Mirroring user's guide*, to complete the remaining steps in this section.

-
14. Complete the appropriate fields as described below.
 - **Original Source**—The name of the source machine where the data originally resided.

- **Restore From**—The name of the target machine that contains the replicated data.
 - **Replication Set**—The name of the replication set to be restored.
 - **Restore To**—The name of the machine where the data will be restored. This may or may not be the same as the original source machine.
15. Identify the correct drive mappings for the data and any other restoration options necessary. For detailed information on the restoration options, see the *HP OpenView Storage Mirroring user's guide*.
 16. Verify that the selections you have made are correct and click **Restore**. The restoration procedure time will vary depending on the amount of data that you have to restore.
 17. After the restoration is complete, start the SQL services on the source machine.
 18. Reestablish the Storage Mirroring replication set and connection.

At this time, your data is restored back to your source machine, the source machine is again the primary server, and, if you selected to continue failover monitoring, the target is available to stand in for the source in the event of a failure.

