

Providing disaster recovery for an HP ProLiant Cluster F200 for MSA1000 with HP OpenView Storage Mirroring



Providing disaster recovery for MSCS	2
Installation notes	2
Installing Storage Mirroring	2
Protecting your data with Storage Mirroring.....	4
Creating a replication set	5
Creating the source connection resource	7
Bring the source connection resource online.....	10
Taking the source connection offline.....	11
Connecting a replication set to a target cluster	11
For more information.....	12

Providing disaster recovery for MSCS

Microsoft Cluster Service (MSCS) leverages the shared storage infrastructure to enable failover of applications and their data. With shared storage, a single copy of the data fails over between the cluster nodes, which ensures that cluster nodes have the latest copy of the data when they host the application. This implementation presents a single point of failure if the data on the shared storage becomes inaccessible or damaged. The ability to have an exact copy of the data mitigates this failure point. HP OpenView Storage Mirroring works within MSCS to enable the creation and maintenance of these copies of the clustered data. Storage Mirroring is powered by NSI Software's Double-Take patented Sequential Transfer Asynchronous Replication technology.

Storage Mirroring replicates from a data source to a target disk system. Data can be replicated in the following cluster configurations:

- A stand-alone source machine connected to a target cluster
- A source cluster connected to a stand-alone target machine
- A source cluster connected to a target cluster

After you decide which configuration suits your needs, you must install Storage Mirroring on any cluster or stand-alone system that will serve as a source or a target. This white paper takes you through the steps to perform cluster-to-cluster replication, which provides you with the maximum uptime of your data.

Installation notes

Before beginning the installation of Storage Mirroring, there are a few things to keep in mind.

- Since Storage Mirroring installs device drivers, it is recommended that you update your Windows Recovery Disk before installing or making changes to your servers. For detailed instructions on creating a recovery disk, see your Microsoft® Windows® reference manuals. Be sure that you select the option to back up the registry when building the repair disks.
- During the installation, it checks to see if MSCS is installed, and if so, installs and registers the cluster resources. If MSCS is not installed when Storage Mirroring is installed, these steps are skipped. If you already have Storage Mirroring installed and then install MSCS, you must perform an Upgrade Install to load the proper dlls for cluster functionality.
- The online user's guide and online help provide additional means of accessing product information. The online manual, dtuser.pdf, is located on the CD in the \docs directory. Review the Preface of the online manual for details on accessing the manual and online help after an installation.
- If you are installing to a drive other than the drive that contains your system TEMP directory, the Microsoft Windows Installer still loads approximately 100 MB of data to the TEMP directory during the installation. If you do not have enough disk space on the drive that contains the TEMP directory, you might need to change where it is located.

Installing Storage Mirroring

Use these instructions if you are installing Storage Mirroring for the first time or if you have completely uninstalled any previous version of Storage Mirroring.

1. Close all open applications.
2. Start the installation program by loading the CD into the local CD-ROM drive. If auto-run is enabled, the program starts automatically. To manually start the program, select **Start>Run**, and specify <cd_drive>:\autorun.exe.

3. The Double-Take installation program starts with an initial screen that enables you to install the software or read associated files. Select **Install HP OpenView Storage Mirroring**.
4. When the Storage Mirroring installation program begins, the Welcome screen is displayed. Click **Next** to continue.
5. Review and accept the Storage Mirroring license agreement to continue with the installation program. Click **Next** to continue.
6. On the Storage Mirroring Setup dialog box, there are two selections available:
 - Client and Server Components—Installs the Storage Mirroring server and client components on the local machine.
 - Client Components Only—Installs the Storage Mirroring client on the local machine.

Select **Client and Server Components**, unless this machine will only be monitoring your connections, and click **Next** to continue.

7. Select the type of Storage Mirroring installation:
 - Complete—All program features will be installed.
 - Custom—You choose which program features will be installed.
8. If you selected the Custom installation, select the components that you want to install and the path where you want the files located. Click **Next** to continue.
9. You will be prompted to enter your user name, organization, and activation code. The code is a 16-character, alpha-numeric activation code that applies the appropriate Storage Mirroring license to your installation. Click **Next** to continue.

NOTE: If you have entered an incorrect or blank activation code, you will be prompted that the code is incorrect and that the Storage Mirroring source and target modules will not load. When prompted to continue, select **Back**, and reenter your activation code.

If you have entered an evaluation code, the expiration date is displayed and you are prompted that the Storage Mirroring source and target modules will not load after that date. When prompted to continue, select **Next** to continue.

10. Review the registration information and confirm that it is accurate by clicking **Next**. If you click **Back**, you will be returned to the registration information screen to correct the information.
11. Storage Mirroring can use disk space to buffer data during periods of high data replication. Specify the size and location of the disk queue. By default, the disk space is set to unlimited, which allows the queue usage to automatically expand whenever the available disk space expands. Click **Next** to continue.
12. The Storage Mirroring security information screen displays next. Review this information, and click **Next** to continue with the installation.
13. If you are satisfied with the selections you have made and are ready to begin copying the Storage Mirroring files, click **Install**.
14. After the files have completed copying, click **Finish** to exit the installation program.

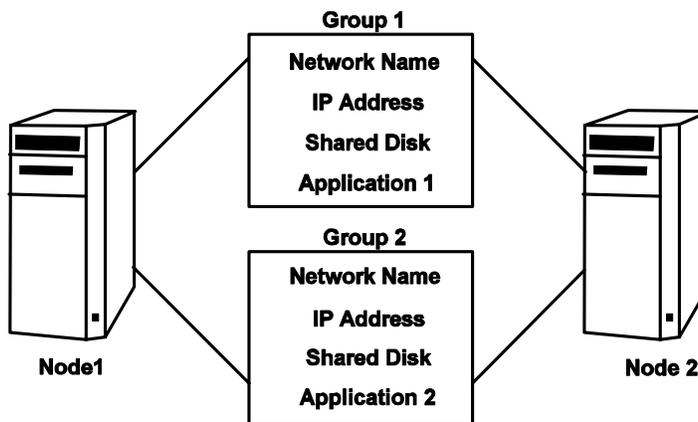
NOTE: If you do not restart your computer when prompted, you must reboot the system before starting Storage Mirroring.

The Storage Mirroring installation is complete.

Protecting your data with Storage Mirroring

The Microsoft Cluster Server configuration contains two machines—Node 1 and Node 2—sharing at least one SCSI disk. Microsoft recommends at least two network interfaces where at least one is dedicated to node-to-node traffic. The cluster contains resources that are services and devices that are managed by the cluster and allowed to move from one machine to the other. The shared SCSI disk is considered a resource as well as drive shares, services, applications, IP addresses, computer names, and so on. Resources are owned by one node of the cluster at a time. If the owning cluster node fails, the second node takes ownership and control. Groups provide an administrative method for collecting multiple resources together for simple management. The groups are actually what moves between the nodes. When a group moves, all of the resources in the group move.

Figure 1.



To provide disaster recovery, you must protect the data that is contained in each group. To provide complete disaster recovery:

1. Using Storage Mirroring Management Console, establish a replication set for each group that contains critical data. You must create the replication set on both machines in the cluster.
2. Using the Microsoft Cluster Administrator, establish a Double-Take Source Connection resource in each group. This resource establishes the Double-Take connection to mirror and replicate data to the target machine.

NOTE: The data will be placed on the same logical drive on the target. Therefore, the target must have the same logical drive configuration as each replication set connected to it.

Creating a replication set

If you are using Storage Mirroring to provide disaster recovery for Microsoft Cluster Server, you must use the following steps when creating a replication set.

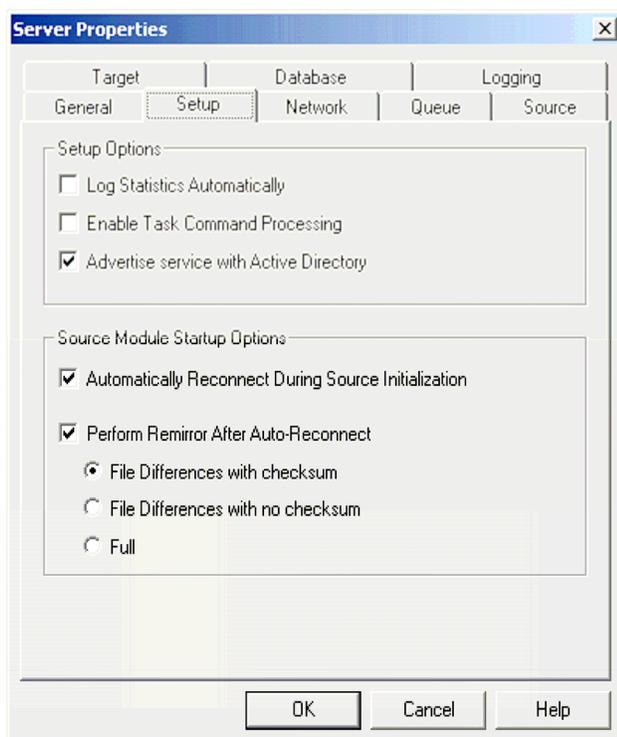
NOTE: If your source is a stand-alone machine and the cluster is the target, you must substitute your stand-alone machine for the node in the following instructions.

1. On the node owning the group you want to protect, select **Start>Programs>HP OpenView Storage Mirroring>Management Console**.
2. Double-click the node name to log on.

NOTE: If your source is a stand-alone machine, skip to step 7.

3. Right-click the node and select **Properties**.
4. Select the **Setup** tab.
5. By default, the Automatically Reconnect During Source Initialization checkbox will be selected. To allow the Double-Take Source Connection resource to fully control the Double-Take connections, disable this option by unselecting the checkbox.

Figure 2.

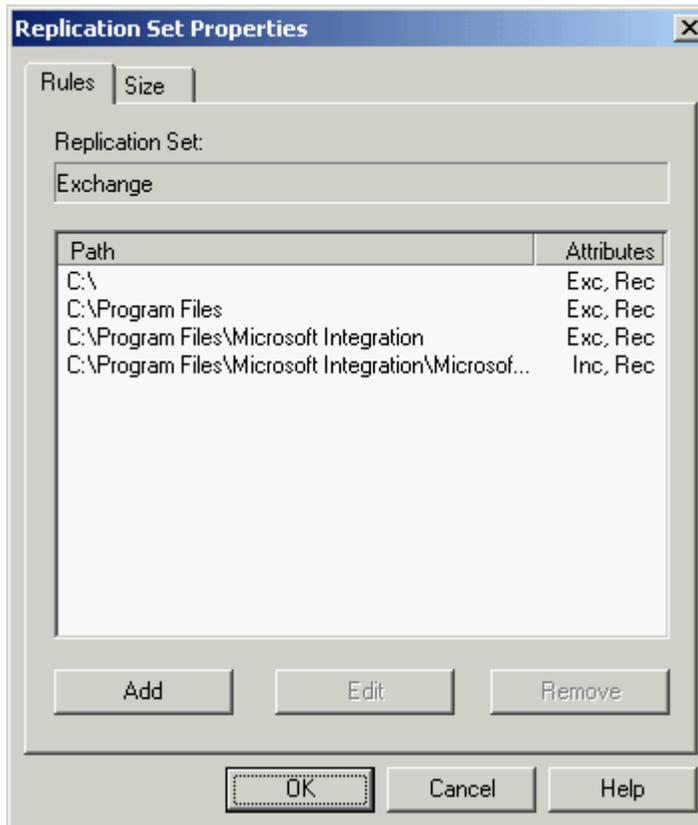


6. Click **OK** to save the changes.
7. Right-click the node and select **New>Replication Set**, and enter the desired name for the replication set.
8. Select the drives and directories on the shared disk that contain the data you want to protect.

NOTE: If there are multiple disks in the same group, you must mark all of the disks. If you have multiple disks in different groups, you must create a unique replication set for each group.

9. Right-click the replication set name, and select **Save** to save the replication set.
NOTE: If your source is a stand-alone machine and the cluster is the target, you can skip the remaining steps in this section and continue with Creating the source connection resource.
10. Right-click the replication set that you just created, and select **Properties**.

Figure 3.



11. Record the exact drive, directory, and attributes for each rule displayed.
12. Double-click the second node in the cluster to log on.
13. Right-click the node, and select **Properties**.
14. Select the **Setup** tab.
15. By default, the Automatically Reconnect During Source Initialization checkbox is selected. Like the first node, disable this option by unselecting the checkbox.
16. Click **OK** to save the changes.
17. Right-click the source, and select **New>Replication Set**. Enter the same name for the replication set as entered on the first node. The replication set name is case-sensitive. Verify that the replication set name on this node is identical, including case-sensitivity, to the replication set on the first node.
18. Right-click the replication set, and select **Properties**.
19. Click **Add**.
20. Specify one of the drive and directory paths that you recorded from the first node. Be sure and mark the correct Include, Exclude, and Recurse subdirectories options that must be applied.

NOTE: Each replication set rule must be identical to the replication set rule on the first node for the disaster recovery process to work correctly.

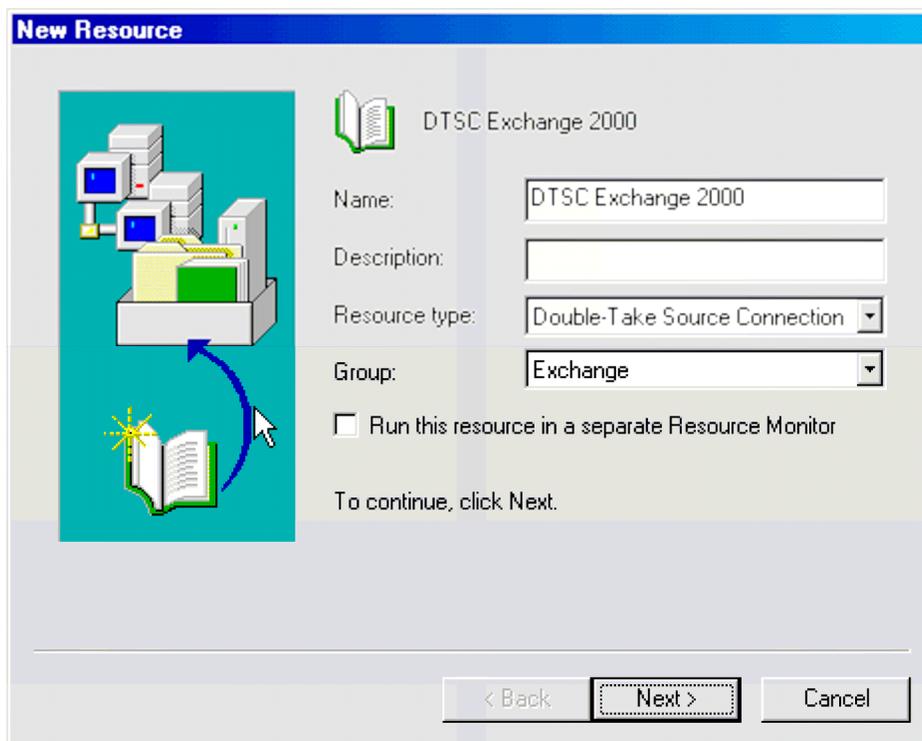
21. Click **OK** to save the replication set rule.
22. Repeat steps 19–21 for each path and directory on the first node.

NOTE: Each drive and directory will appear in the Replication Set Properties even though the second node may not have access to these locations right now. This is not a problem.
23. Right-click the replication set name, and select **Save** to save the replication set.

Creating the source connection resource

1. Select **Start>Programs>Administrative Tools>Cluster Administrator**.
2. Select the group for which you want to provide disaster recovery.
3. Right-click that group, and select **New>Resource**.

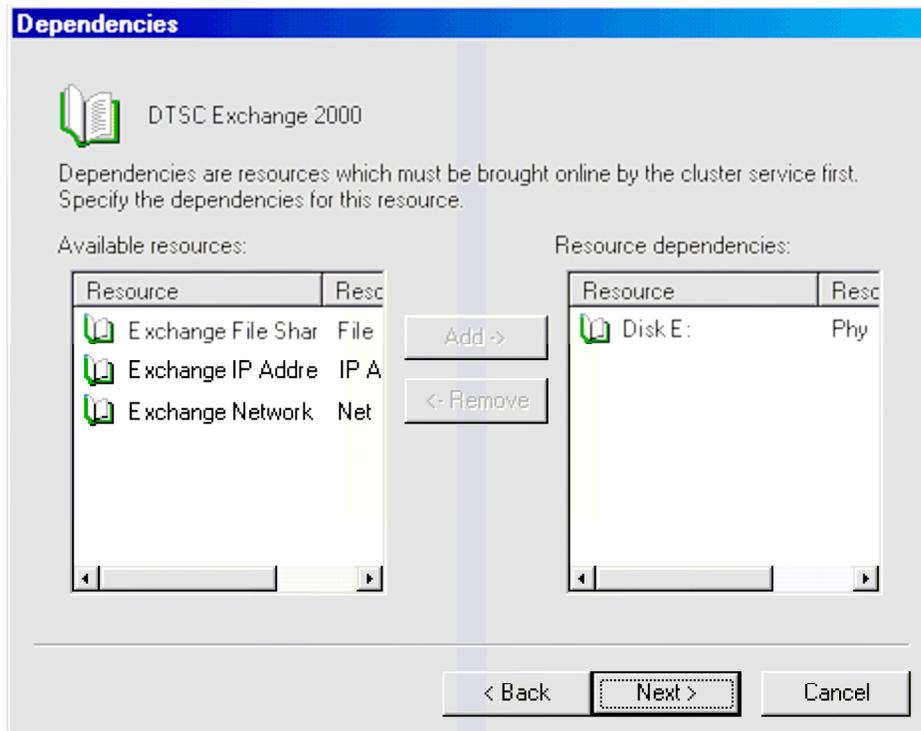
Figure 4.



4. Specify the following fields on the New Resource dialog box:
 - **Name**—Specify the name of this resource. You may want the name to reflect the type of data that is being protected.
 - **Description**—Specify a description for this resource.
 - **Resource type**—Specify Double-Take Source Connection.
 - **Group**—Verify the group from where you originated the new resource is the correct group.
5. Click **Next** to continue.
6. The default on the Possible Owners dialog box includes both nodes as possible owners. This default does not need to be changed. Click **Next** to continue.

7. The resource is dependent on the disks that contain the data that was included in your replication set. Verify that the resource dependencies include your physical disk resources, and click **Next** to continue.

Figure 5.



8. Specify the Double-Take Source Connection Parameters:

- Replication Set—The name of the replication set that you are using. This name is case-sensitive and must be typed exactly as it was entered in the Management Console.
- Double-Take Target—The name of the machine that will receive the replicated data.

NOTE: If you are using an MSCS as your target, enter the virtual IP address in the Double-Take Target properties.

- Allow connection scripts to interact with the desktop—Select this checkbox if you want to display the connection information in a command prompt dialog box.

NOTE: If the script is set to interact with the desktop, the results are displayed on the owning node. The resource remains at the Online pending status until the script is manually acknowledged by a user. Therefore, this option should not be a permanent setting, but should be used if you are having problems establishing a connection.

- Logon to Target as—Allows you to specify an ID, password, and domain to be used when logging on to the target.

NOTE: The user must be a member of the Double-Take Admin security group on all nodes of the cluster running a Double-Take Source Connection resource.

Figure 6.

Double-Take Source Connection Parameters

DTSC Exchange 2000

Specify the replication set and Double-Take target to connect. To setup intra-cluster replication, specify a cluster node as the Double-Take target.

Replication Set: (case-sensitive) Exchange

Double-Take Target: JERSEY

Allow connection scripts to interact with the desktop

Logon to Target as:

Username: administrator

Password: *****

Domain: domain.com

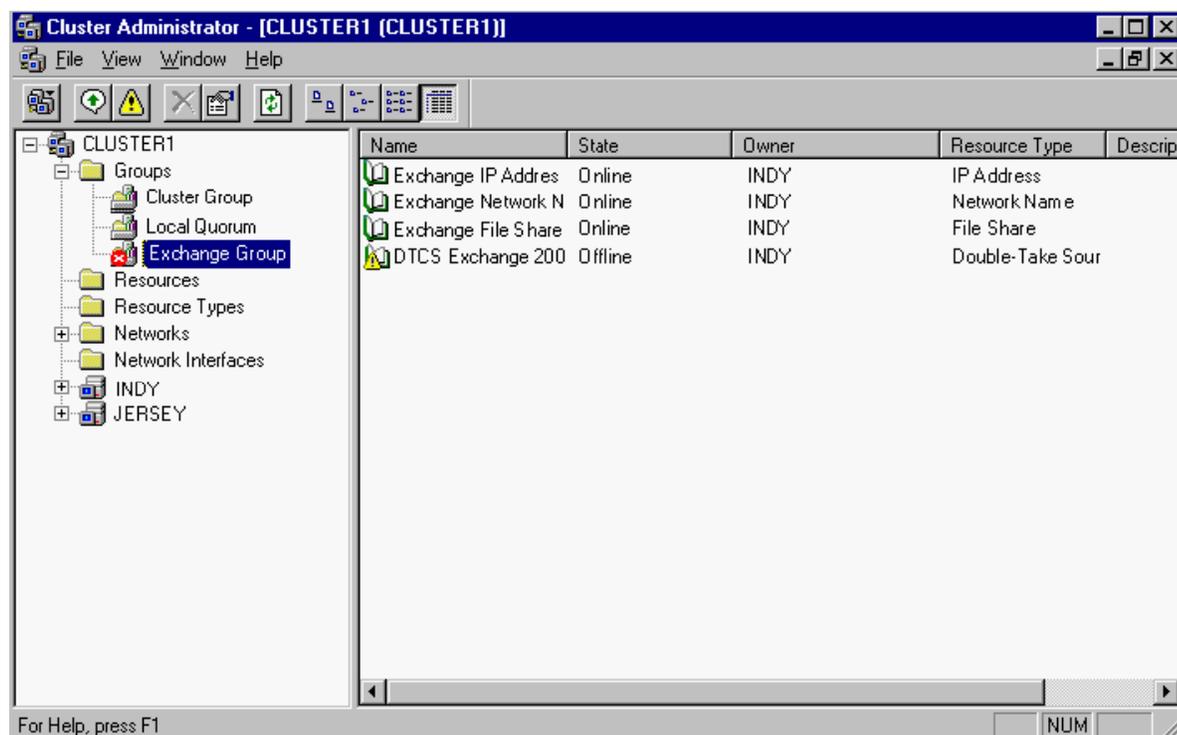
< Back Finish Cancel

9. Click **Finish** to complete the creation of the Double-Take Source Connection resource.

Bring the source connection resource online

Your new resource is displayed in an offline state in Cluster Administrator similar to the following example.

Figure 7.



Notice that the offline resource in the right pane of the Cluster Administrator contains a warning sign icon that appears as a yellow triangle with an exclamation point. The group in the left pane of Cluster Administrator that contains the resource has a warning sign icon that appears as a red circle with an X.

Bring the resource online by right-clicking the resource and selecting **Bring Online**.

The warning signs are removed from the resource and the group. The State column indicates the resource is online and functioning properly. The data is now being replicated from the source to the target.

NOTE: If the State column indicates failed, the connection has not executed properly. When the resource has failed, the Storage Mirroring connection is not active and your data is not being protected in the event of a disaster.

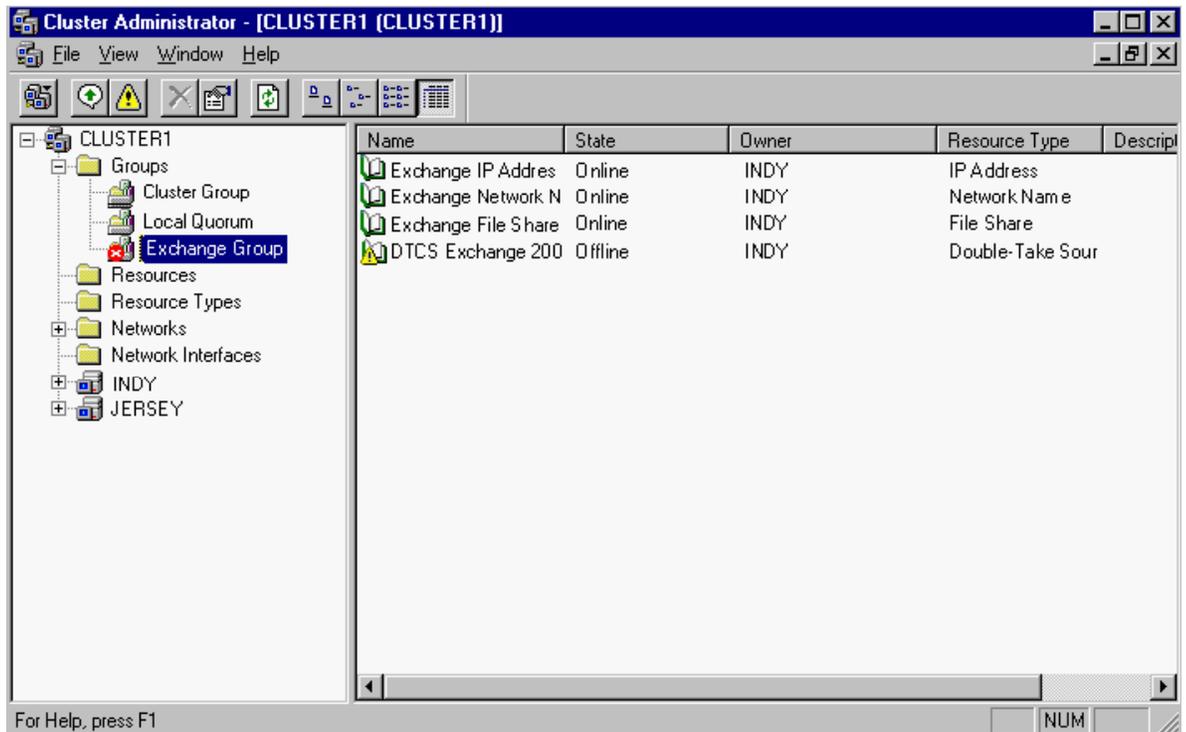
To troubleshoot a failed Double-Take Source Connection resource, right-click the resource, and select **Properties**. Select the **Parameters** tab. Select the **Allow connection script to interact with desktop** checkbox, and click **OK**. When you attempt to bring the resource online, the results from the connection script are displayed on screen on the node running the script allowing you to see what is causing the failure.

When using the Allow connection script to interact with desktop option with the Double-Take Source Connection resource, the resource reports a pending state until the command prompt is closed.

Taking the source connection offline

If you must terminate the Double-Take Source Connection resource, right-click the resource, and select **Take Offline**. The resource state changes to offline in Cluster Administrator.

Figure 8.



If the Double-Take Source Connection resource is offline, data is not being replicated and is not protected in the event of a disaster.

Connecting a replication set to a target cluster

If you are replicating data from a stand-alone machine to a cluster that is acting as the Double-Take target, select **Start>Programs>HP OpenView Storage Mirroring>Text Client**. The connection must be made through the Text Client by using the connect command with the <target_machine> option set to the IP address assigned to the group that contains the shared disk to which you are replicating.

After the connection is established, the data is replicated to the shared disk of the cluster. If the node ownership changes, the connection moves to the new owner since the connection is to the virtual server's IP address of the cluster.

For more information

<http://h18000.www1.hp.com/solutions/enterprise/highavailability/index.html>

© 2004 Hewlett-Packard Development Company, L.P. The information contained herein 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 and Windows are U.S. registered trademarks of Microsoft Corporation.

5982-5666EN, 05/2004

