



High Availability for Clustered Exchange Servers Using Double-Take and the Double-Take Exchange Failover Utility



High Availability for Clustered Exchange Servers Using Double-Take and the Double-Take Exchange Failover Utility
Revision 1.3.0 published October 2004

Double-Take, GeoCluster and NSI are registered trademarks of NSI Software, Inc. Balance is a trademark of NSI Software, Inc. Microsoft, Windows, and the Windows logo are trademarks, or registered trademarks of Microsoft Corporation in the United States and/or other countries. All other trademarks are the property of their respective companies.

© 1996-2004 NSI Software, Inc.

Double-Take Support for Application Failover

The Double-Take[®] file system replication process is application independent and replicates any file system changes (including permissions and attributes) written to NTFS, FAT or FAT32 file systems by any application or process, subject to specific exceptions called out in the *User's Guide* or *readme.txt* file. Maintaining point-in-time consistent file system replicas and providing server monitoring and automatic or manual failover of the server name and IP address are the primary functions of the Double-Take software and we offer support to qualified customers should these functions fail to operate in accordance with our published documentation, regardless of what application or process is manipulating the data.

NSI[®] Software may provide application notes and other documents that provide implementation guidelines on how to use Double-Take functions and replicas to manually or automatically failover or recover many popular third party applications and a general process to accomplish failover or recovery of many other third party applications. While these steps are believed to be accurate for the specific configuration, Double-Take version, and application versions originally tested, due to the number of possible configurations and variables, NSI Software can only test selected combinations and may provide only limited support for the operation and configuration of third party applications or the behavior of those applications before, during, or after failover, in its discretion. In cases where NSI Software has no direct access to or experience with a particular application or configuration, NSI Software support may also be limited to only the actual replication of the file system data and failover (name and IP address) of the server.

For assistance in validating, implementing or troubleshooting these or other possible configurations with third party applications, NSI Software and its partners may offer professional services on a fee basis to apply best practices for assisting with third party applications to recover automatically or manually using replicated data. This, and any other, application note is provided solely for the convenience of our customers and is not intended to bind NSI Software to any obligation. Although we try to provide quality information, NSI Software makes no claims, promises or guarantees about the accuracy, completeness, or adequacy of the information contained in this document.

Table of Contents

1 Introduction	1
2 Requirements	1
3 Backing Up Your Environment	1
4 Configuring the Source Cluster	2
4.1 Installing Exchange	2
4.2 Creating the Exchange virtual server	2
4.2.1 Creating the Exchange group	2
4.2.2 Creating the Exchange IP Address resource	3
4.2.3 Creating the Exchange Network Name resource	4
4.2.4 Creating the Physical Disk resource	6
4.2.5 Creating the Exchange System Attendant Resource	7
4.3 Installing and configuring Double-Take	8
4.4 Configuring the Double-Take replication set	8
4.5 Configuring the Double-Take Source Connection resource	9
5 Configuring the Target Cluster	11
5.1 Installing Exchange	11
5.2 Creating the Exchange virtual server	11
5.3 Installing and configuring Double-Take	11
6 Bringing the Components Online and Beginning Monitoring	12
7 Dealing With a Failure	13
8 Failing Back to Your Source Cluster	15
Appendix I: Configuring Additional Exchfailover.exe Utility Options	17

1. Introduction

Microsoft® Exchange Server is a messaging and collaboration server for the most demanding business needs. Its scalability, performance, and enhanced security make Exchange an ideal messaging foundation for enterprise networks. Double-Take provides real-time enterprise data protection and replication for cluster environments. Double-Take can be used to provide high availability for your Exchange server.

This document describes the steps necessary to configure Double-Take to provide high availability and disaster recovery for two clusters running Windows® 200x, MSCS, and Microsoft Exchange Server version 2000 or 2003.

In addition to the Exchange Information Stores (mailboxes and public folders), there are many important aspects of an Exchange server configuration that are required for Exchange functionality following the loss of a production Exchange Server. It is important to be aware of the overall production Exchange server configuration, and to configure the target identically. Some configuration aspects fall outside the scope of this document such as the configuration of any Exchange Connectors, Built-In Instant Messaging, Newsgroups, Bridgehead Servers, and so on. These issues need to be addressed exclusively by the Exchange administrator.

To complete these instructions, you will install Microsoft Exchange Server and Double-Take, and configure Double-Take for replication and failover. Due to the complexities of these applications, this document is intended for network administrators with experience installing, configuring, and maintaining network applications, including Double-Take, MSCS, and Microsoft Exchange Server.

2. Requirements

Each node in each cluster must meet the following system requirements

- ◆ One of the following operating systems:
 - ◆ Microsoft Windows 2000 Advanced Server with Service Pack 2 or later
 - ◆ Microsoft Windows Server 2003 Enterprise Edition

NOTE: All nodes must both be running the same operating system.

- ◆ One of the following Microsoft Exchange Server versions:
 - ◆ Microsoft Exchange Server 2000 Enterprise Edition with Service Pack 3 or higher
 - ◆ Microsoft Exchange Server 2003 Enterprise Edition

NOTE: The Exchange version must be the same as the operating system version. (Exchange 2000 is not supported in a Windows 2003 environment nor is Exchange 2003 in a Windows 2000 environment.)

- ◆ A licensed copy of Double-Take version 4.3.4 or later
- ◆ A copy of the NSI Software ExchFailover.exe utility version 1.24. This utility is available on the Application Notes page of the NSI Software support web site, where you downloaded this file from.
- ◆ Each node must be member servers in the same domain.

3. Backing Up Your Environment

Before beginning these procedures, make sure you have a current backup of all nodes. Also make sure you have a complete backup of Active Directory.

4. Configuring the Source Cluster

For the source cluster you will need to complete the following:

- 4.1 Installing Exchange
- 4.2 Creating the Exchange virtual server
- 4.3 Installing and configuring Double-Take
- 4.4 Configuring the Double-Take replication set
- 4.5 Configuring the Double-Take Source Connection resource

NOTE: If Exchange is already installed on your source cluster, you can skip sections 4.1 and 4.2. But you must make sure you know the configuration information (storage groups, databases, paths, and file names) so that the target cluster can be configured identically.

If Double-Take is already installed, review section 4.3 to make sure your configuration matches the required configuration for this Exchange solution.

4.1. Installing Exchange

Download the appropriate Microsoft document that is appropriate for your Exchange version and service pack level to prepare the forest and domain and install Exchange on each node of the source cluster. These documents can be found on the Microsoft web site. Record critical Exchange configuration information from the source cluster, including storage groups, databases, paths, and file names so that this information can be used when installing Exchange on the target.

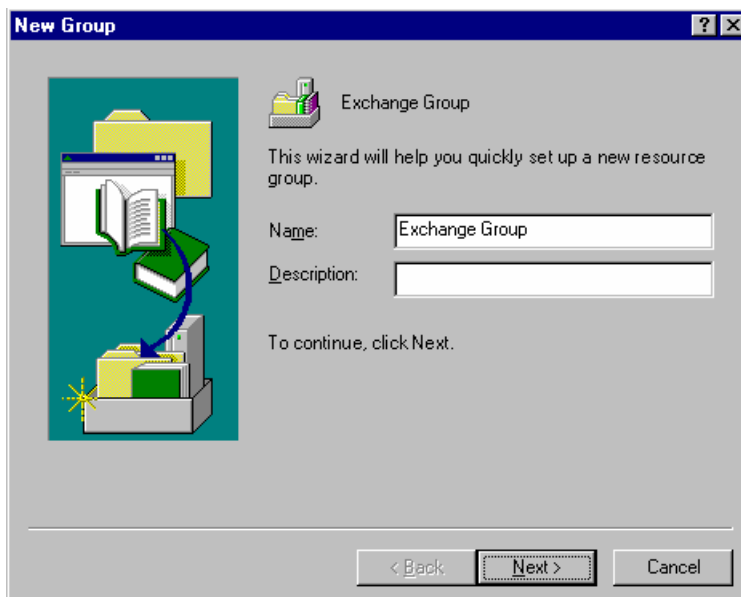
4.2. Creating the Exchange virtual server

To create the Exchange virtual server, you will be creating an Exchange group and adding the following resources:

- ◆ IP address
- ◆ Network name
- ◆ Physical Disk Resource
- ◆ Exchange System Attendant

4.2.1. Creating the Exchange group

1. Right-click the **Groups** folder on the left pane of the Cluster Administrator and select **New, Group**.




2. Specify the **Name** and **Description** for the Exchange group and click **Next** to continue.
3. No **Preferred Owners** are required. Click **Finish** to complete the creation of the new group.


NOTE: You will be notified that the group was created successfully. Click **OK** to acknowledge the message and return to the Cluster Administrator main screen.

4.2.2. Creating the Exchange IP Address resource

1. Right-click the Exchange group that you just created and select **New, Resource**.
2. Specify the following fields on the New Resource dialog box:



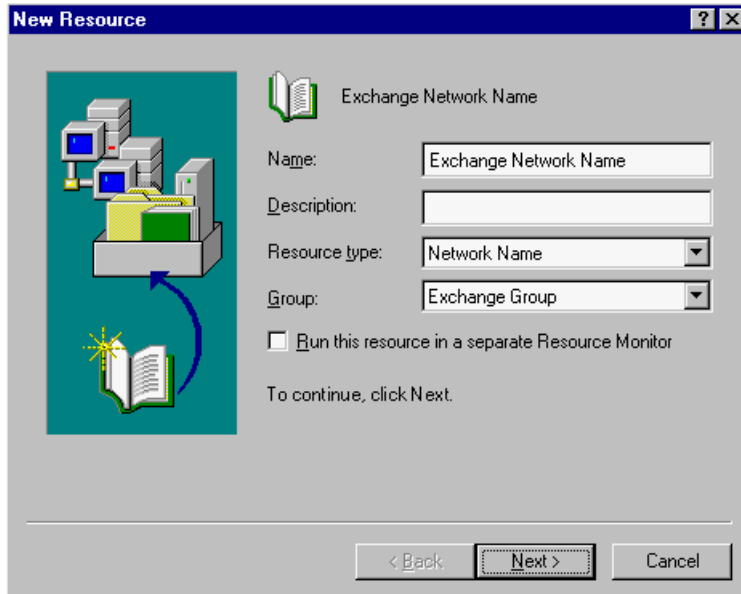
- ◆ **Name**—Specify a name that identifies this resource as the IP address for the Exchange group. This name must be unique within the cluster.
 - ◆ **Description**—You can optionally add a more detailed description for this resource.
 - ◆ **Resource type**—Specify **IP Address**.
 - ◆ **Group**—The Exchange group should be selected. If it is not, select the correct group name.
3. Click **Next** to continue.
 4. The default **Possible Owners** does not need to be modified. Click **Next** to continue.
 5. There are no **Dependencies** required. Click **Next** to continue.
 6. Specify the following fields on the TCP/IP Address Parameters dialog box:



-
- ◆ **Address**—Enter the IP address that will be assigned to the Exchange virtual server.
 - ◆ **Subnet mask**—Enter the subnet mask associated with the IP address you just entered.
 - ◆ **Network to use**—If you have more than one route for network traffic defined, specify the network that this IP address will use.
7. Click **Finish** to complete the creation of the IP address resource.
 8. Right-click the IP address resource and select **Bring Online**.

4.2.3. Creating the Exchange Network Name resource

1. Right-click the Exchange group and select **New, Resource**.
2. Specify the following fields on the New Resource dialog box:



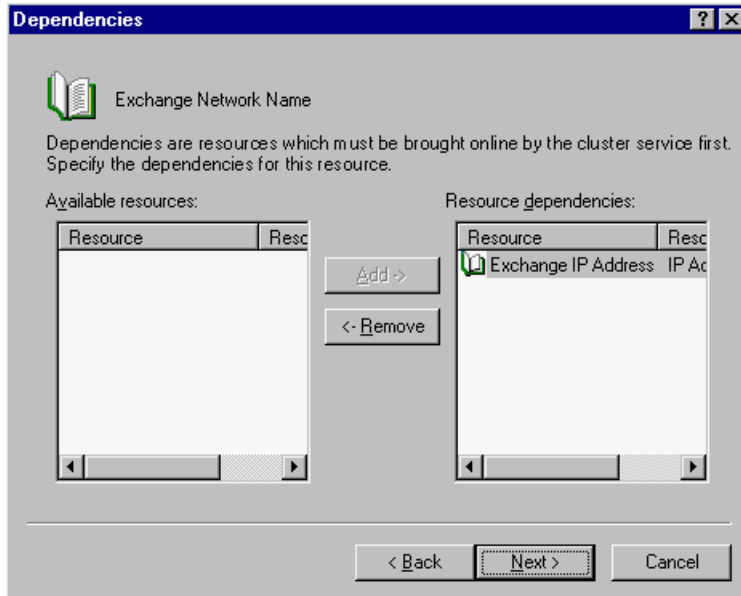
The screenshot shows the 'New Resource' dialog box with the following fields and options:

- Name:** Exchange Network Name
- Description:** (empty text box)
- Resource type:** Network Name (dropdown menu)
- Group:** Exchange Group (dropdown menu)
- Run this resource in a separate Resource Monitor
- To continue, click Next.

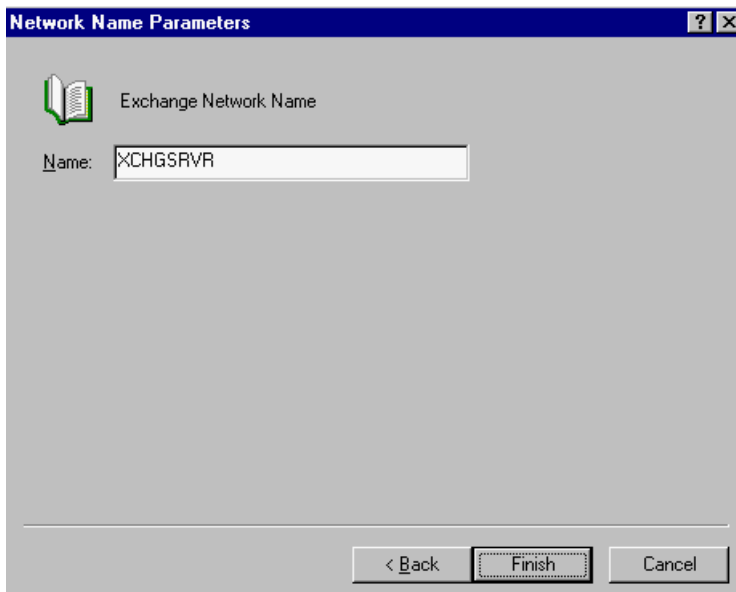
At the bottom of the dialog box are three buttons: '< Back', 'Next >', and 'Cancel'. The 'Next >' button is highlighted with a dashed border.

- ◆ **Name**—Specify a name that identifies this resource as the virtual server name for the Exchange group. This name must be unique within the cluster.
 - ◆ **Description**—You can optionally add a more detailed description for this resource.
 - ◆ **Resource type**—Specify **Network Name**.
 - ◆ **Group**—The Exchange group should be selected. If it is not, select the correct group name.
3. Click **Next** to continue.
 4. The default **Possible Owners** does not need to be modified. Click **Next** to continue.

- An IP address must be present in order for a network name to be assigned. Therefore, move the IP address resource that you just created to the **Resource dependencies** list and click **Next** to continue.




- Specify the Network Name Parameters by entering the virtual name of the Exchange server. This is the name that clients will look for on the network.




- Click **Finish** to complete the creation of the Network Name resource.
- Right-click the Network Name resource and select **Bring Online**.

4.2.4. Creating the Physical Disk resource

1. Right-click the group and select **New, Resource**.
2. Specify the following fields on the New Resource dialog box:



- ◆ **Name**—Specify a name that identifies the disk drive associated with the virtual server. This name must be unique within the cluster.
 - ◆ **Description**—You can optionally add a more detailed description for this resource.
 - ◆ **Resource type**—Specify **Physical Disk**.
 - ◆ **Group**—The application group name should be selected. If it is not, select the correct group name.
3. Click **Next** to continue.
 4. Verify that both nodes appear as **Possible Owners** and click **Next** to continue.
 5. No resources are required as dependencies. Click **Next** to continue.
 6. Specify the disk drive associated with the physical disk resource.

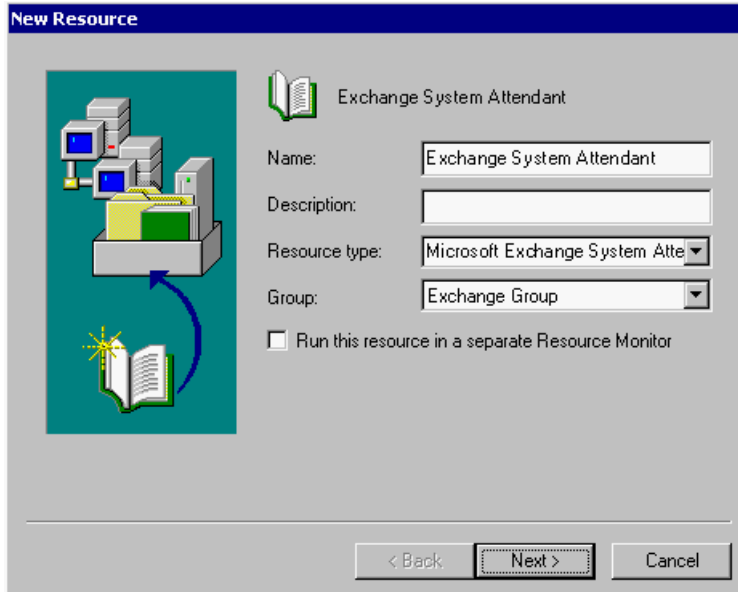


7. Click **Finish** to complete the creation of the Physical Disk resource.
8. Right-click the physical disk resource and select **Bring Online**.

9. Some Exchange implementations separate the databases and log files on to different disk resources for improved disk I/O. If this is the case in your environment, repeat steps 1-8 and create additional physical disk resources for each disk.

4.2.5. Creating the Exchange System Attendant Resource

1. In the Cluster Administrator, right-click the Exchange group and select **New, Resource**.
2. Specify the following fields on the New Resource dialog box:



The screenshot shows the 'New Resource' dialog box in Cluster Administrator. The title bar reads 'New Resource'. On the left, there is a graphic of a server rack with a blue arrow pointing to a book icon. The main area is titled 'Exchange System Attendant'. It contains the following fields and controls:

- Name:** A text box containing 'Exchange System Attendant'.
- Description:** An empty text box.
- Resource type:** A dropdown menu showing 'Microsoft Exchange System Atte'.
- Group:** A dropdown menu showing 'Exchange Group'.
- Run this resource in a separate Resource Monitor

At the bottom, there are three buttons: '< Back', 'Next >' (which is highlighted with a dashed border), and 'Cancel'.

- ◆ **Name**—Specify a name that identifies the Exchange System Attendant.
 - ◆ **Description**—You can optionally add a more detailed description for this resource.
 - ◆ **Resource type**—Specify **Microsoft Exchange System Attendant**.
 - ◆ **Group**—The Exchange group should be selected. If it is not, select the correct group name.
3. Click **Next** to continue.
 4. The Exchange System Attendant resource is dependent on the network name and the physical disk resource. Move the Exchange Server Name and the physical disk resource that you just created to the **Resource dependencies** list and click **Add** to continue.
 5. Verify the location of the Exchange data files, for example E:\Exchsrvr, and the click **Finish**.
 6. Right-click the Exchange System Attendant resource and select **Bring Online**.

4.3. Installing and configuring Double-Take

1. Install Double-Take on each node of the source cluster using the installation defaults.
2. Disable the standard connection controls so that the Double-Take Source Connection resource, that you will be configuring later, can control the Double-Take connections.
 - a. Start the Double-Take Management Console by selecting **Start, Programs, Double-Take, Management Console**.
 - a. Double-click on the first node on the left pane of the Management Console to login.
 - b. Right-click the first node of the cluster and select **Properties**.
 - c. Select the **Setup** tab.
 - d. By default, the **Automatically Reconnect During Source Initialization** check box will be selected. Disable this option by clearing the check box.
 - e. Click **OK** to save the changes.
3. Enable block checksum comparisons on differences mirrors.
 - a. Right-click the first node of the cluster and select **Properties**.
 - b. Select the **Source** tab.
 - c. By default, the **Block Checksum All Files on a Difference Mirror** check box will not be selected. Enable this option by marking the check box.
 - d. Click **OK** to save the changes.
4. Repeat steps 2 and 3 on the second node of the cluster.

4.4. Configuring the Double-Take replication set

In order for the clusters to be synchronized, the data that is changed on the source cluster must be replicated to the target cluster. Double-Take handles this task by establishing a replication set which identifies the data that is changing.

1. Assuming the Double-Take Management Console is still open from the last section, right-click the node owning the group you wish to protect and select **New, Replication Set**.
2. Enter a name for the replication set and press **Enter**.
3. On the right pane of the Management Console, select the directory and files in the tree that are associated with the group and virtual server you are protecting.
4. Right-click the replication set name and select **Save**.
5. Right-click the replication set that you just created and select **Properties**.
6. Record the exact drive and directories of each path displayed in the Replication Set Properties table below. Place a check mark or X in the Include, Exclude, and Recurse Sub-directories columns to identify which parameters apply to the specified path.

Replication Set Properties

Drive and Directories	Include	Exclude	Recurse Sub-directories

7. Double-click on the second node on the left pane of the Management Console to login.

8. Right-click the node and select **New, Replication Set**.
9. Enter the exact, case-sensitive name for the replication set as specified on the first node and press **Enter**.
10. Right-click the replication set that you just created and select **Properties**.
11. Click **Add**.
12. Specify one of the drive and directory paths that you recorded in the table *Replication Set Properties* on page 8. Be sure and mark the correct **Include**, **Exclude** and **Recurse sub-directories** options that need to be applied.

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

13. Click **OK** to save the replication set rule.
14. Repeat steps 13-15 for each path and directory on the first node.

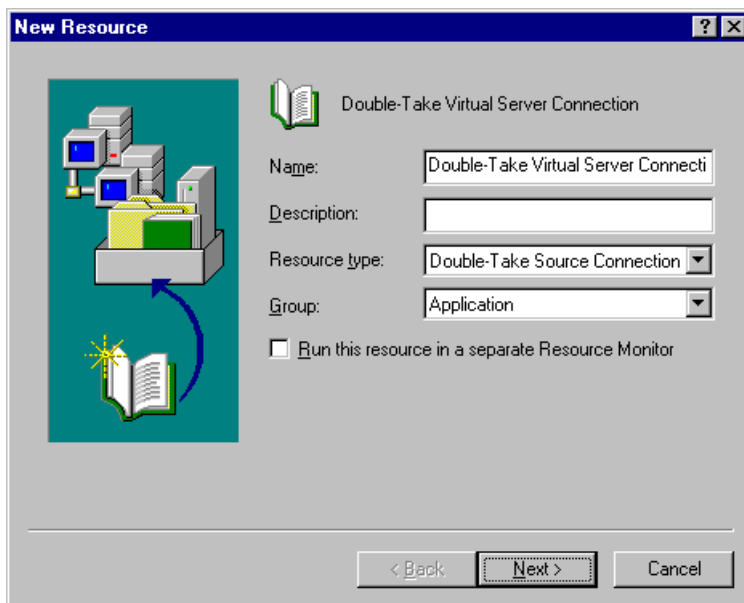
NOTE: Each drive and directory will appear in the Replication Set Properties even though the second node may not have access to these locations right now. That is not a problem.

15. Right-click the replication set name and select **Save**.

4.5. Configuring the Double-Take Source Connection resource

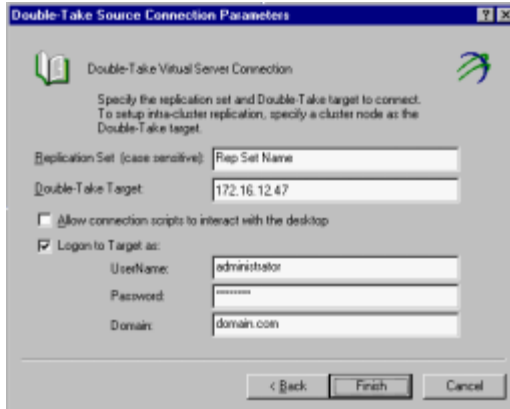
The Double-Take Source Connection resource controls the Double-Take connections. You need to configure this resource through the Cluster Administrator.

1. Select **Start, Programs, Administrative Tools, Cluster Administrator**.
2. Right-click the application group you were working with earlier and select **New, Resource**.
3. Specify the following fields on the New Resource dialog box:



- ◆ **Name**—Specify a name that indicates this is the Double-Take virtual server connection.
 - ◆ **Description**—You can optionally add a more detailed description for this resource.
 - ◆ **Resource type**—Specify **Double-Take Source Connection**.
 - ◆ **Group**—The application group name should be selected. If it is not, select the correct group name.
4. Click **Next** to continue.
 5. Verify that both nodes appear as **Possible Owners** and click **Next** to continue.
 6. To keep the Double-Take Source Connection resource from coming online before the physical disk, make this resource dependent on the physical disk resource just created. Click **Next** to continue.

7. Specify the following on the Double-Take Source Connection Parameters dialog box:



- ◆ **Replication Set**—Specify the name of the Double-Take replication set. This name is case-sensitive and should be the same name as specified in [Configuring the Double-Take replication set](#) on page 8.
- ◆ **Double-Take Target**—In the next section, you will be creating two IP Address resources on the target cluster. One will be offline and will be used for client access in the event of a failure and the target cluster becomes the primary cluster. The second IP Address resource will be for Double-Take communications between the two clusters. In the **Double-Take Target** field of the Double-Take Source Connection Parameters dialog box, specify the unique IP address that you will be using for Double-Take communications between the two clusters.

The other options on the Double-Take Source Connection Parameters dialog box are optional. See the *User's Guide* for additional details.

8. Click **Finish** to complete the creation of the Double-Take Source Connection resource.
9. You will need to modify one of the scripts that controls the Double-Take Source Connection resource to enable orphan file removal.
 - a. Open the file `online.dtcl`, located in the directory where you installed Double-Take, with a text editor.
 - b. Locate the following section in the file.

```
if ($exitcode = 0) then
    WRITE "Starting block checksum mirror";
    $mirror = MIRROR START $conid DIFFERENT,CHECKSUM;
```

- c. Add a space and the keyword `ORPHANS` (before the semicolon) to the last line in this section so that it is identical the following command:
`$mirror = MIRROR START $conid DIFFERENT,CHECKSUM ORPHANS;`
 - d. Save and close the file.
 - e. Repeat steps a-d on the second node of the cluster.

5. Configuring the Target Cluster

Like the source configuration, you will be installing Exchange on the target, creating the Exchange virtual server, and configuring Double-Take, but you won't be bringing Exchange online.

5.1. Installing Exchange

Install Exchange on the target cluster with the same configuration matching the source cluster including storage groups, databases, paths, and file names.

5.2. Creating the Exchange virtual server

1. Repeat the steps in [4.2 Creating the Exchange virtual server](#) on page 2 to create another Exchange virtual server. The resources in this group should be unique from those on the source cluster.

NOTE: Make sure that none of the resources are online after you create them.

2. The IP address created on the target in step 1 above is offline while the target cluster is the Double-Take target. If a failure should occur and the target cluster becomes the primary cluster that the clients are accessing, this IP address will be brought online during failover. Therefore, the target cluster has to have another IP address that Double-Take can use for its communications. In the resource group that you created in step 1 above, create another IP Address resource specific for Double-Take communications. This should be the same IP address that you specified for the **Double-Take Target** on the Double-Take Source Connection Parameters dialog box in [Configuring the Double-Take Source Connection resource](#) on page 9.

5.3. Installing and configuring Double-Take

1. Install Double-Take on each node of the target cluster using the installation defaults.
2. Disable the Double-Take standard connection controls so that the Double-Take Source Connection resource can control the Double-Take connections.
 - a. Start the Double-Take Management Console by selecting **Start, Programs, Double-Take, Management Console**.
 - b. Right-click the first node of the target cluster and select **Properties**.
 - c. Select the **Setup** tab.
 - d. By default, the **Automatically Reconnect During Source Initialization** check box will be selected. Disable this option by clearing the check box.
 - e. Click **OK** to save the changes.
3. Enable block checksum comparisons on differences mirrors.
 - a. Right-click the first node of the cluster and select **Properties**.
 - b. Select the **Source** tab.
 - c. By default, the **Block Checksum All Files on a Difference Mirror** check box will not be selected. Enable this option by marking the check box.
 - d. Click **OK** to save the changes.
4. Repeat steps 2 and 3 on the second node of the target cluster.

6. Bringing the Components Online and Beginning Monitoring

The final step is to bring the components online and for the target cluster to begin monitoring the source cluster for a failure.

1. Open the source cluster in Cluster Administrator (**Start, Programs, Administrative Tools, Cluster Administrator**) and right-click the Exchange group and select **Bring Online**. Verify that all of the resources in the group are brought online successfully.
2. Configure the target cluster to monitor the source.
 - a. On a node of the target cluster, select **Start, Program Files, Double-Take, Failover Control Center**.
 - b. Select the node you are running on from the **Target Machine** list box.
 - c. Click **Add Monitor**.
 - d. You need to enter an IP address to monitor instead of a source machine name, click **Custom**.
 - e. Specify the IP address and network name for the Exchange virtual server and click **OK**.
 - f. Select the IP address to be monitored by marking the check box to the left the address in the **Names to Monitor** tree.
 - g. Clear all of the settings in the **Items to Failover** section.
 - h. Clear the **Use .SHR Share Mapping File (If Available)** option.
 - i. Click **OK** to save and close the monitor information.
3. Repeat step 2 on the second node of the target cluster.

Double-Take is now replicating from the source to the target cluster and the target cluster is now monitoring the IP address of the source cluster. In the event that both nodes of the source cluster fail, Double-Take will bring the resources online on the target cluster. For details on monitoring failover, see the *User's Guide*.

7. Dealing With a Failure

If a failure occurs and the Failover Control Center Time to Fail counter reaches zero (0), a dialog box will appear in the Failover Control Center requiring user intervention to initiate failover. (If the Failover Control Center is not open when the failure occurs, the dialog box will appear the next time the Failover Control Center is opened. See the *User's Guide* for details on monitoring a failure.) Acknowledge the manual intervention prompt to start failover.

You must complete the following steps to bring the target Exchange group online.

NOTE: You can automate these steps by using a script. The necessary commands are shown in the note below.

1. The `exchfailover.exe` utility must be executed with the same permissions that the MS Exchange SA service uses, which is by default System permission. You have two alternatives for using these permissions.

- a. Add the user you are currently logged in as to the Exchange Domain Servers security group. This group contains the appropriate System permissions. Open a command prompt window on the target cluster node that owns the Exchange group after you have added yourself to the Exchange Domain Servers group.
- b. Open a command prompt on the target node and enter the following command:

```
at \\node_name hh:mm /interactive "cmd"
```

where `source_name` is the name of the source and `hh:mm` is the current time plus one minute, formatted using a 24 hour clock. For example, if the current time is 3:00pm, then you would enter 15:01. The command would be `at \\node_name 15:01 interactive "cmd"`. In one minute, another command prompt window will open on the node console (not through Terminal Services) using the appropriate System permissions.

2. In the command prompt window (the new window if you used the `at` command above), change to the Double-Take directory and execute the following command:

```
exchfailover -setup -failover -s source_Exchange_virtual_name -t  
target_Exchange_virtual_name
```

where `source_Exchange_virtual_name` is the name of the source Exchange virtual server and `target_Exchange_virtual_name` is the name of the target Exchange virtual server.

Do not close the command window at this time.

3. Bring the target Exchange group online.
4. After the Exchange group is online, you will need to run the Exchange failover utility again to prepare the databases for use. In the same command prompt window used in step b, execute the following command:

```
exchfailover -failover -s source_Exchange_virtual_name -t target_Exchange_virtual_name
```

where `source_Exchange_virtual_name` is the name of the source Exchange virtual server and `target_Exchange_virtual_name` is the name of the target Exchange virtual server.

NOTE: If your source server is the routing master, you will need to modify the security settings so that the routing master role can be failed over. If your source server is not the routing master, skip this step and continue with step 3.

You have two options available for modifying the security settings for the routing master role.

- ◆ The first option grants the `exchfailover.exe` utility the permission to perform the task for you. Edit the `postover.bat` file that you just created and add the `-u username:password` switch as outlined in the table [ExchFailover Utility Command Syntax](#) on page 19. Specify the Exchange administrator account information.
 - ◆ The second option allows you to set the security setting manually, thus not requiring the `-u` switch in the failover and failback scripts. For this second option, you need to modify the properties of the Active Directory entry noted below by adding the System account with write rights to the entry as well as enabling the option **This object and all child objects to a particular attribute** on the Advanced tab of the Active Directory entry. The Active Directory entry that must be edited is `CN=Routing Groups, CN=First Administrative Group, CN=Administrative Groups, CN=Exchange_Organziation_Name, CN=Microsoft Exchange, CN=Services, CN=Configuration, DC=Domain_name, DC=com`. You will have to use `ADSIedit` from the Windows Support Tools to make this change. See your Windows reference guide for more details.
-

After the command completes, clients can connect through Outlook or Outlook web access to receive their e-mail.

NOTE: You can automate these steps by creating a script containing the following commands. The commands are case-sensitive and you will need to substitute the name of your source virtual server, target virtual server, and the Exchange group.

```
exchfailover.exe -setup -failover -s source_Exchange_virtual_name -t target_Exchange_virtual_name
cluster group "Group_Name" /ONLINE
"c:\program files\doubletake\dtcl.exe" wait 25000
exchfailover.exe -failover -s source_Exchange_virtual_name -t target_Exchange_virtual_name
```

8. Failing Back to Your Source Cluster

If your source cluster experiences a failure causing both nodes to fail, such as a power outage, your target cluster will stand in for the source while you resolve the issues. During this downtime, users may be updating data on the target cluster. When your source cluster is ready to come back online, the data is no longer current and must be updated with the new data on the target cluster.

1. Bring one or both of the production cluster nodes online, but do not allow the Exchange group to start.

NOTE: If necessary, you may have to disconnect one cluster node from the production network, bring it online, take the Exchange resource group offline, and then reconnect the node to the network.

2. Using the instructions in section [4.4 Configuring the Double-Take replication set](#) on page 8, create a replication set on the target cluster that includes all of the volumes contained in the target Exchange resource group.
3. Using the instructions in section [4.5 Configuring the Double-Take Source Connection resource](#) on page 9, create the Double-Take Source Connection resource to the source cluster.
4. Bring the Double-Take Source Connection resource online so that mirroring to the source cluster begins.
5. Monitor the mirroring process to ensure that it completes with no errors. Users may continue to access Exchange on the target cluster server.
6. After the mirror completes, schedule downtime to convert operations back to the source cluster. At the scheduled time, bring the Exchange resources offline on the target cluster, leaving the Double-Take Source Connection resource online.
7. Monitor the Queued Replication Bytes field in the Double-Take Management Console to ensure that all data has been replicated to the production cluster. Review the Double-Take logs on the active node of the production cluster to ensure that there are no write errors.
8. Once all of the data has been replicated, take the Double-Take Source Connection resource offline.
9. Complete the following steps to bring the source Exchange group online.

NOTE: You can automate these steps by using a script. The necessary commands are shown in the note below.

- a. The `exchfailover.exe` utility must be executed with the same permissions that the MS Exchange SA service uses, which is by default System permission. You have two alternatives for using these permissions.
 1. Add the user you are currently logged in as to the Exchange Domain Servers security group. This group contains the appropriate System permissions. Open a command prompt window on the source cluster node that owns the Exchange group after you have added yourself to the Exchange Domain Servers group.
 2. Open a command prompt on the source node and enter the following command:

```
at \\node_name hh:mm /interactive "cmd"
```

where `source_name` is the name of the source and `hh:mm` is the current time plus one minute, formatted using a 24 hour clock. For example, if the current time is 3:00pm, then you would enter 15:01. The command would be `at \\node_name 15:01 interactive "cmd"`. In one minute, another command prompt window will open on the node console (not through Terminal Services) using the appropriate System permissions.
- b. In the command prompt window (the new window if you used the `at` command above), change to the Double-Take directory and execute the following command:

```
exchfailover -setup -failback -s source_Exchange_virtual_name -t target_Exchange_virtual_name
```

where `source_Exchange_virtual_name` is the name of the source Exchange virtual server and `target_Exchange_virtual_name` is the name of the target Exchange virtual server.

Do not close the command window at this time.
- c. Bring the source Exchange group online.
- d. After the Exchange group is online, you will need to run the Exchange failover utility again to prepare the databases for use. In the same command prompt window used in step b, execute the following command:

```
exchfailover -failback -s source_Exchange_virtual_name -t target_Exchange_virtual_name
```

where `source_Exchange_virtual_name` is the name of the source Exchange virtual server and `target_Exchange_virtual_name` is the name of the target Exchange virtual server.

After the command completes, clients can connect through Outlook or Outlook web access to receive their e-mail.

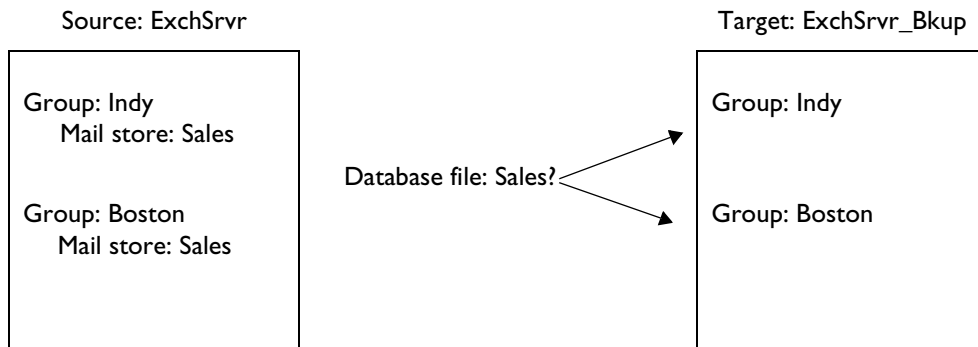
NOTE: You can automate these steps by creating a script containing the following commands. The commands are case-sensitive and you will need to substitute the name of your source virtual server, target virtual server, and the Exchange group.

```
exchfailover.exe -setup -failback -s source_Exchange_virtual_name -t target_Exchange_virtual_name
cluster group "Group_Name" /ONLINE
"c:\program files\doubletake\dtcl.exe" wait 25000
exchfailover.exe -failback -s source_Exchange_virtual_name -t target_Exchange_virtual_name
```

Appendix I: Configuring Additional Exchfailover.exe Utility Options


In order for a mail store (and its users) to be failed over (or failed back), a mail store on the source must be paired to a mail store on the target. In order to be a valid pair, the database filename (excluding path information) of these two stores must match. Exchfailover.exe uses two methods to make these mail store pairs. The simplest (default) method requires that the database filenames be unique, that it only occur once on the source, and that it only occur once on the target. If your environment uses the same store name in different groups or if you need to rename stores or groups on the target during failover, you will need to add additional options to the exchfailover.exe utility used in the postover.bat script.

For example, a server called ExchSrvr contains two mail groups Indy and Boston. Each group contains a mail store called Sales. In its simplest form, the exchfailover.exe utility would not know which group to associate the Sales mail store with since it is based on the database file name.



To resolve this issue, you can direct the groups and mail stores to meet your environment needs. The `-r` option in the exchfailover.exe utility is a pairing rule. It allows you to specify how the groups and mail stores on the source will be paired on the target.

By itself, the `-r` option will create a one-to-one mapping from the source to the target. For example, the command `exchfailover.exe -failover -s ExchSrvr -t ExchSrvr_Bkup -r` would automatically create a one-to-one mapping on the target.



You can be more specific with the `-r` option and direct the source groups to specific group names on the target. For example, the command

```
exchfailover.exe -failover -s ExchSrvr -t ExchSrvr_Bkup -r Indy:Indy_Bkup -r Boston:Boston_Bkup
```


will pair the mail stores from the source Indy group in the group Indy_Bkup on the target. And the mail stores from the source Boston group will be paired in the group Boston_Bkup on the target.



If needed, you can be the most specific with the `-r` option by specifying both the group and mail store names. For example, if you need to direct the group and mail store names on the target, the command

```
exchfailover.exe -failover -s ExchSrvr -t ExchSrvr_Bkup -r Indy, Sales:Indy_Bkup, Sales -r Boston, Sales:Boston_Bkup, Sales
```

will pair the mail store Sales in the Indy_Bkup group from the Sales mail store from the Indy group on the source. It will also pair the mail store Sales in the Boston_Bkup group from the Sales mail store from the Boston group on the source.



There are several other options available in the `exchfailover.exe` utility. These options and the full command syntax are included in the table on the next page.

ExchFailover Utility Command Syntax

Command EXCHFALLOVER

Description Used in script files to failover Exchange data

Syntax EXCHFALLOVER **-FAILOVER** | **-FAILBACK** **-s** <source> **-t** <target> [**-l** <log_filename>] [**-norus**] [**-nospn**] [**-o** <options_filename>] [**-r** [<source_group>][,<source_mail_store>][:[<target_group>][,<target_mail_store>]]] [**-SETUP**] [**-test**] [**-u** <username>:<password>] [**-?[?]**]

- Options**
- ◆ **FAILOVER**—The Exchange data will be moved from the source to the target during failover
 - ◆ **FAILBACK**—The Exchange data will be moved from the target to the source during failback
 - ◆ **s** *source*—The name of the original source server
 - ◆ **t** *target*—The name of the original target server
 - ◆ **l** *log_filename*—The name of the optional log file name. By default, the log file is ExchFailover.log and is stored in the directory containing the exchfailover.exe file. If this name is changed, the DTInfo utility will not be able to locate this file which could impede assistance through Technical Support.
 - ◆ **norus**—Do not change the Recipient Update Service
 - ◆ **nospn**—Do not change the Service Principle Name
 - ◆ **o** *options_filename*—Allows you to pass in a file containing the options for the exchfailover utility
 - ◆ **r**—By itself, this option creates a one-to-one mapping of the groups and mail stores from the source to the target
 - ◆ **r** *source_group:target_group*—The r option with the group names will direct the source group name specified to the target group name specified
 - ◆ **r** *source_group, source_mail_store:target_group, source_mail_store*—The r option with all of the r options will direct the source group name and mail store specified to the target group name and mail store specified
 - ◆ **SETUP**—Allows you to set the overwrite database on restore flag without completing user moves or RUS and folder updates. If the -setup switch is not supplied, the utility still sets the overwrite database on restore flag, but the other work is performed also.
 - ◆ **test**—Test mode that does not change the Exchange configuration
 - ◆ **u** *username:password*—A user with Active Directory permissions
 - ◆ **?**—Displays the syntax of the exchfailover.exe utility
 - ◆ **??**—Displays the syntax of the exchfailover.exe utility along with brief descriptions of each option

Examples

- ◆ `exchfailover -failover -s Indy -t ExchSrvr_Bkup`
- ◆ `exchfailover -failover -s Indy -t ExchSrvr_Bkup -r`
- ◆ `exchfailover -failover -s Indy -t ExchSrvr_Bkup -r Sales:Indy_Sales`
- ◆ `exchfailover -failover -s Indy -t ExchSrvr_Bkup -r Sales, Inside:Indy_Sales, Inside -r Sales, Outside:Indy_Sales, Outside`
- ◆ `exchfailover -failover -s Indy -t ExchSrvr_Bkup -r Sales:Indy_Sales -norus -u administrator:password`
- ◆ `exchfailover -failover -s Indy -t ExchSrvr_Bkup -o options_file.txt`

- Notes**
- ◆ When using the -failback option, the source related options pertain to your original source or what will become the new source, if the original source had to be replaced The target related options pertain to the target, that is currently standing in for the source.
 - ◆ The password specified with the -u option is the only case-sensitive option in this command.