Double-Take Availability User's Guide

Double-Take, Balance, Double-Take Availability, Double-Take Backup, Double-Take Cargo, Double-Take Flex, Double-Take for Hyper-V, Double-Take for Linux, Double-Take Move, Double-Take ShadowCaster, Double-Take for Virtual Systems, GeoCluster, Livewire, netBoot/i, NSI, sanFly, TimeData, TimeSpring, winBoot/i and associated logos are registered trademarks or trademarks of Double-Take Software, Inc. and/or its affiliates and subsidiaries in the United States and/or other countries. Microsoft, Hyper-V, 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.

© 1996-2010 Double-Take Software, Inc. All rights reserved.

DTA-521-UG

3/22/2010

Table of Contents

Double-Take Availability overview

Double-Take Availability ensures the availability of critical workloads. Using real-time replication and failover, you can protect data, individual applications, entire servers, or virtual machines.

Identify your critical workload on your production server, known as the source, and replicate the workload to a backup server, known as the target. The target server, on a local network or at a remote site, stores the copy of the workload from the source. Double-Take Availability monitors any changes to the source workload and sends the changes to the copy stored on the target server. By replicating only the file changes rather than copying an entire file, Double-Take Availability allows you to more efficiently use resources.

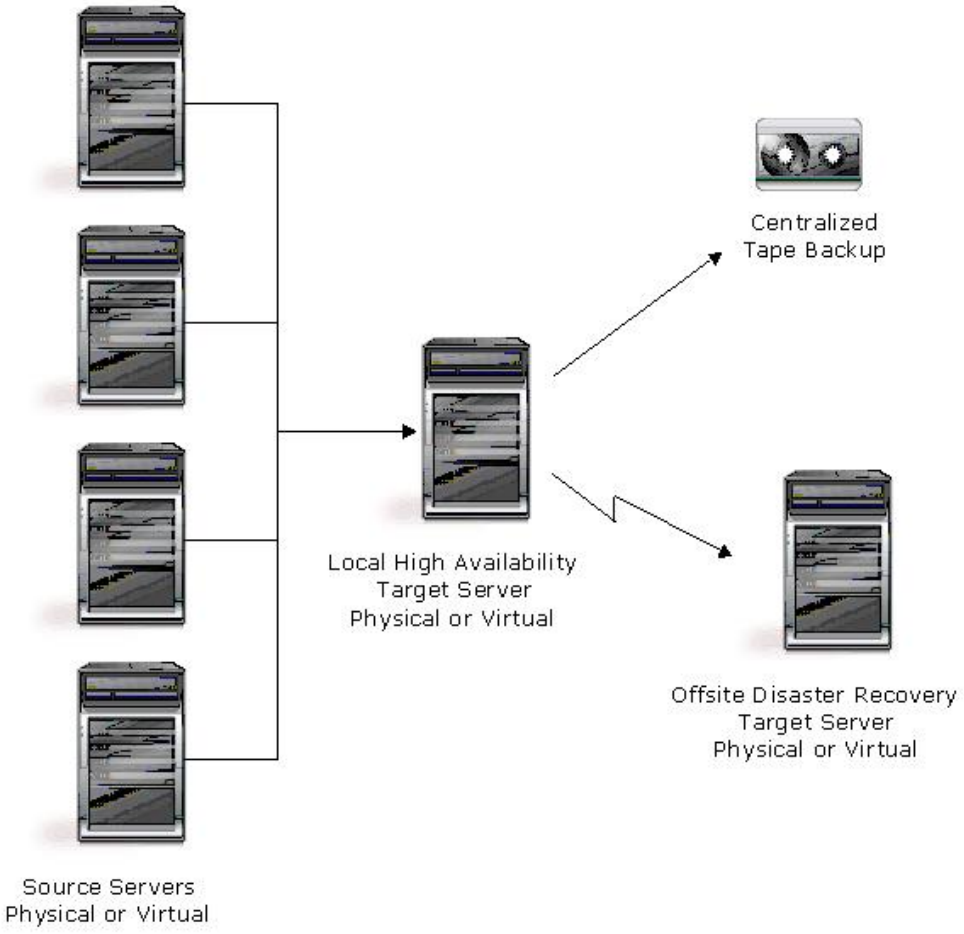

Core operations

Double-Take Availability performs four basic types of operations.

- Mirroring—The initial copy or subsequent resynchronization of selected data
- Replication—The on-going capture of byte-level file changes
- Failure monitoring and failover—The ability to monitor and stand-in for a server, in the event of a failure
- Restoration—A mirror of selected data from the target back to the source

Mirroring

Mirroring is the process of transmitting user-specified data from the source to the target so that an identical copy of data exists on the target. When Double-Take Availability initially performs mirroring, it copies all of the selected data, including file attributes and permissions. Mirroring creates a foundation upon which Double-Take Availability can efficiently update the target server by replicating only file changes.

If subsequent mirroring operations are necessary, Double-Take Availability can mirror specific files or blocks of changed data within files. By mirroring only files that have changed, network administrators can expedite the mirroring of data on the source and target servers.

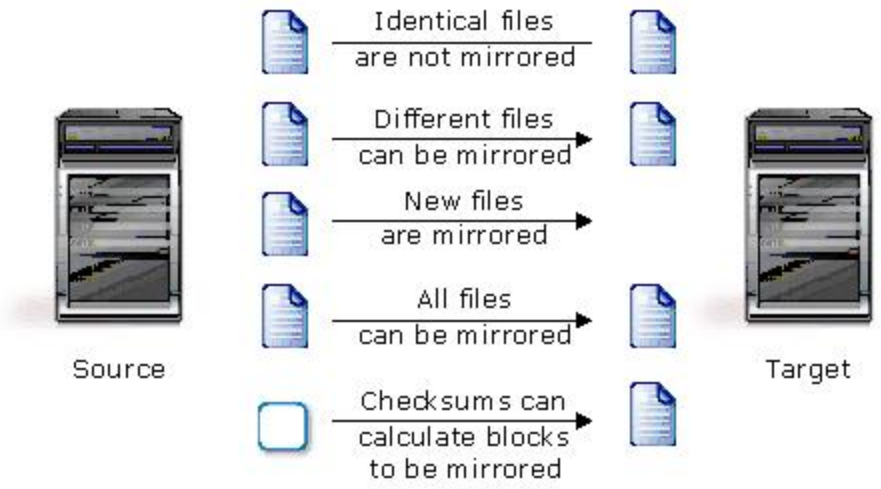

Mirroring has a defined end point - when all of the selected files from the source have been transmitted to the target. When a mirror is complete, the target contains a copy of the source files at that point in time.

Replication

Replication is the real-time transmission of file changes. Unlike other related technologies, which are based on a disk driver or a specific application, the Double-Take Availability replication process operates at the file system level and is able to track file changes independently from the file's related application. In terms of network resources and time, replicating changes is a more efficient method of maintaining a real-time copy of data than copying an entire file that has changed.

After a source and target have been connected through Double-Take Availability, file system changes from the user-defined data set are tracked. Double-Take Availability immediately transmits these file changes to the target server. This real-time replication keeps the data on the target up-to-date with the source and provides high availability and disaster recovery with minimal data loss.

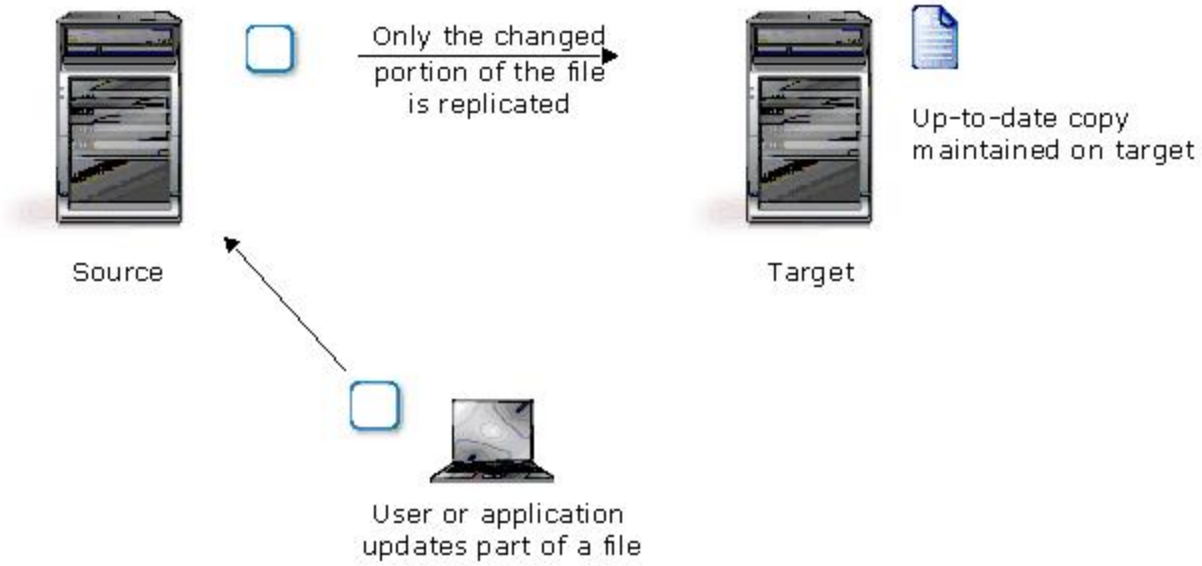

Unlike mirroring which is complete when all of the files have been transmitted to the target, replication continuously captures the changes as they are written to the source. Replication keeps the target up-to-date and synchronized with the source.

Failure monitoring and failover

Failover is the process in which a target stands in for a failed source. As a result, user and application requests that are directed to the failed source are routed to the target.

Double-Take Availability monitors the source status by tracking network requests and responses exchanged between the source and target. When a monitored source misses a user-defined number of requests, Double-Take Availability assumes that the server has failed. Double-Take Availability then prompts the network administrator to initiate failover, or, if configured, it occurs automatically.

The failover target assumes the network identity of the failed source. When the target assumes the identity of the source, user and application requests destined for the source server or its IP address(es) are routed to the target.

When partnered with the Double-Take Availability data replication capabilities, failover routes user and application requests with minimal disruption and little or no data loss. In some cases, failover may be used without data replication to ensure high availability on a server that only provides processing services, such as a web server.

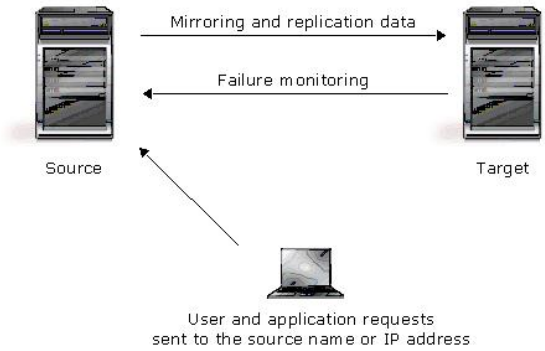

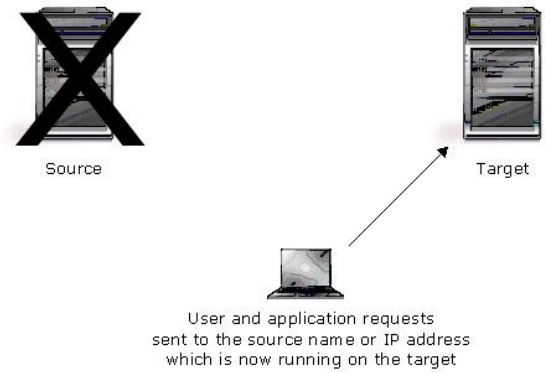

Restoration

Restoration provides an easy method for copying replicated data from the target back to its original location on the source. The process only requires you to select the source, target, and the appropriate replication set. There is no need to select files or to remember where the data came from on the source since that information is maintained by Double-Take Availability.

Restoration can be used if the source data is lost due to a disk crash or when the most up-to-date data exists on the target due to failover. At the time of a source server failure, your Double-Take Availability target will contain the same data as your Double-Take Availability source. If you are using the Double-Take Availability failover capabilities, users can continue updating data on the target server while the problems on the source are resolved. Because of the continued updates on the target, when the source server is ready to come back online, the two servers will no longer contain the same data. Restoration is the process of copying the up-to-date data from the target back to the original source or a new source.

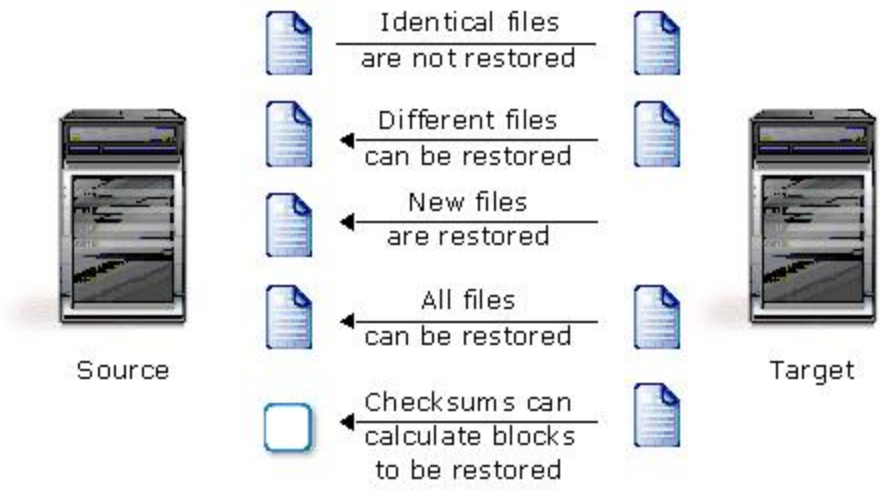

When a restoration is complete, the source and target are again synchronized. Replication continues from the target to the source, keeping the two servers synchronized, until you disconnect the restoration connection.

Double-Take Availability workloads

Building on Double-Take Availability core operations, you can protect specific workloads to meet your protection goals.

- Full-server workloads—You can protect an entire server, including the data and system state, which is the server's configured operating system and applications. In the event of a failure, the target becomes the source.
- Application workloads—You can protect applications running on your source including Exchange, SQL, SharePoint, BlackBerry, or a Windows file server.
- Virtual workloads—You can protect virtual servers in the following configurations.
  - You can protect an entire physical or virtual server to an automatically provisioned (created) virtual server on a Hyper-V or ESX server. If you are protecting a virtual server, you are protecting the data within the guest operating system.
  - You can protect a Hyper-V virtual server to a Hyper-V virtual server. In this case, you are protecting the host-level files (.vhd files), making them highly available on another Hyper-V server.
  - You can protect an ESX virtual server to an ESX virtual server. In this case, you are protecting the host-level files (.vmdk files), making them highly available on another ESX server.
- Cluster workloads—You can protect two types of clusters.
  - You can protect a standard cluster where a single copy of data resides on a SCSI disk shared between cluster nodes.
  - You can protect a GeoCluster that eliminates the single point of failure of a shared disk by replicating data between volumes.

Full-server workloads

Full-server workload protection provides high availability for an entire server, including the system state, which is the server's configured operating system and applications. You identify your source, which is the server you want to protect, and your target, which is the server that will stand-in for the source in the event the source fails. Once the two servers are selected and their configurations validated, Double-Take Availability monitors the source for a failure. When it fails, Double-Take Availability allows the target to stand-in for the source by rebooting and applying the source, including its system state, on the target. After the reboot, the target becomes the source, and the target no longer exists.

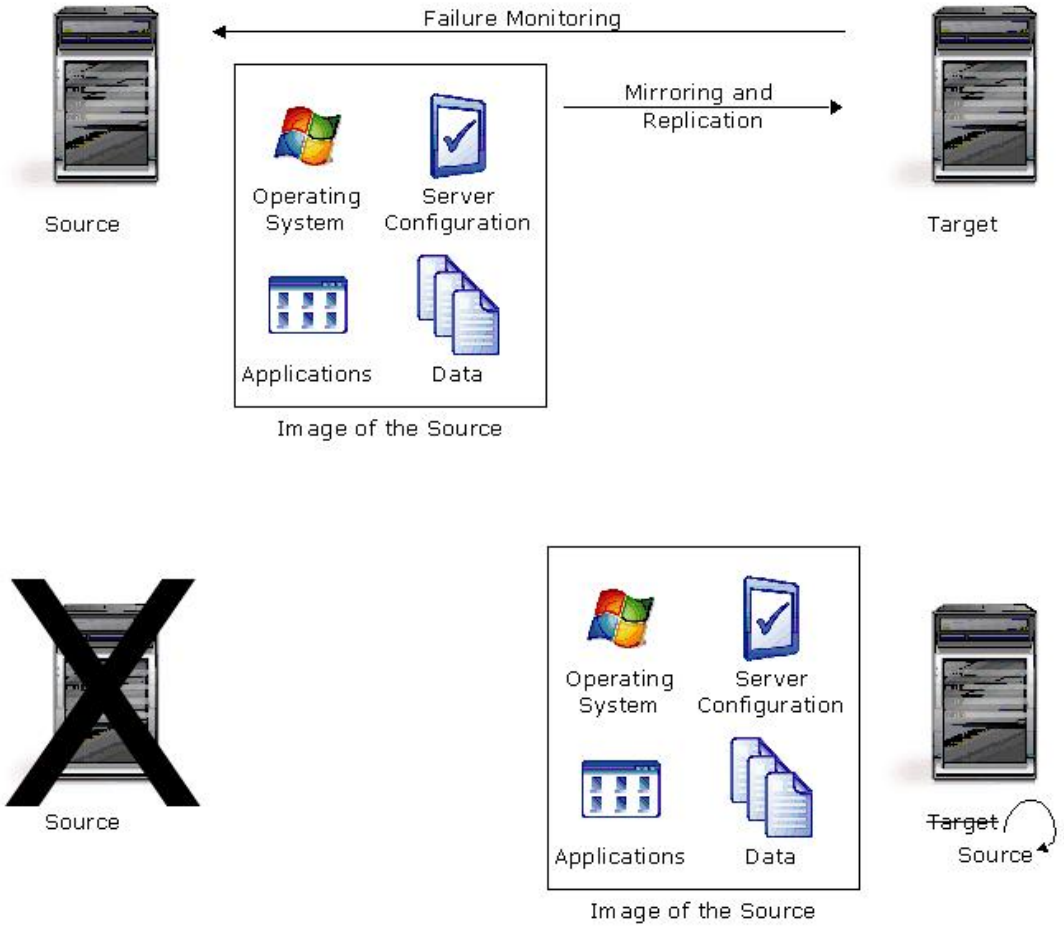

Application workloads

Application workload protection provides high availability for Exchange, SQL, SharePoint, BlackBerry, and a Windows file server. You identify your source, which is the server running the application, and your target, which is the server that will stand-in for the source in the event the source fails. Double-Take Availability will gather information from your environment (Active Directory, DNS, and so on) about the application being protected and automatically protect the application. Double-Take Availability monitors the source server or the application services for a failure. When it fails, Double-Take Availability allows the target to stand-in for the source. Your end-users can continue to access the application running on the target, which is standing in for the source.

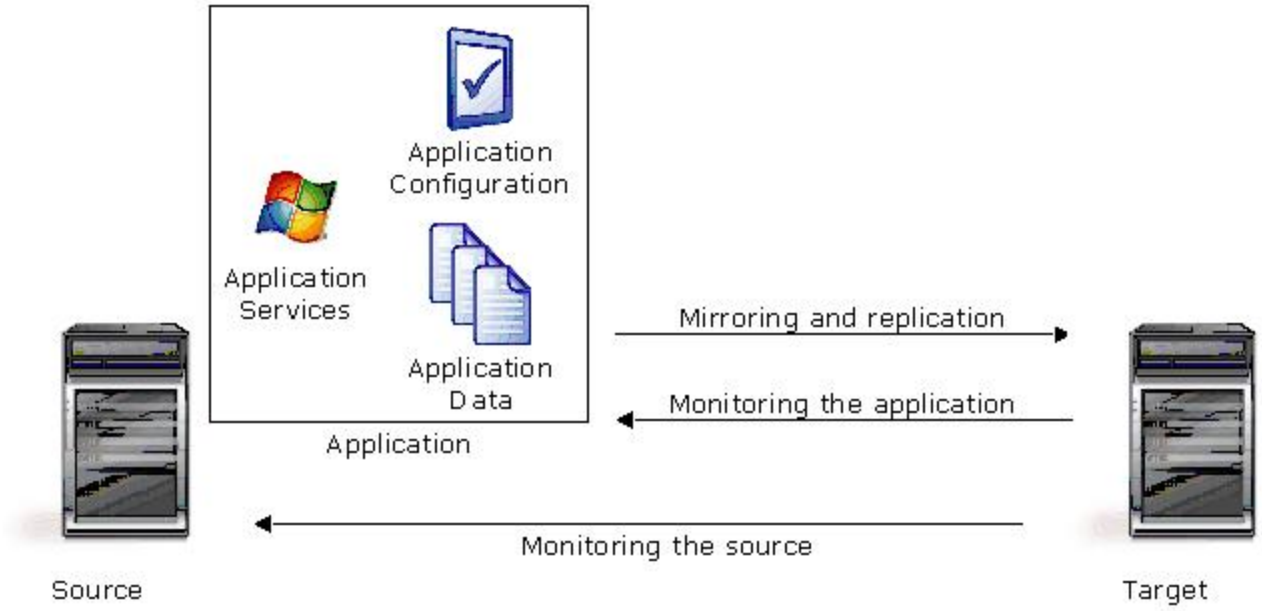

Virtual workloads

Virtual workload protection provides high availability to Hyper-V or ESX virtual servers. You identify your source, which is the server you want to protect. Your source can be a physical server, a virtual machine where you want to protect the data within the guest operating system, or a virtual machine where you want to protect the host-level virtual disk files (.vhd or .vmdk files). Your target is a Hyper-V or ESX server that will host a virtual machine that is a replica of the source. Double-Take Availability monitors the source for a failure. In the event of a source failure, the replica virtual machine on the target can stand-in allowing end-users to continue accessing data and/or applications.

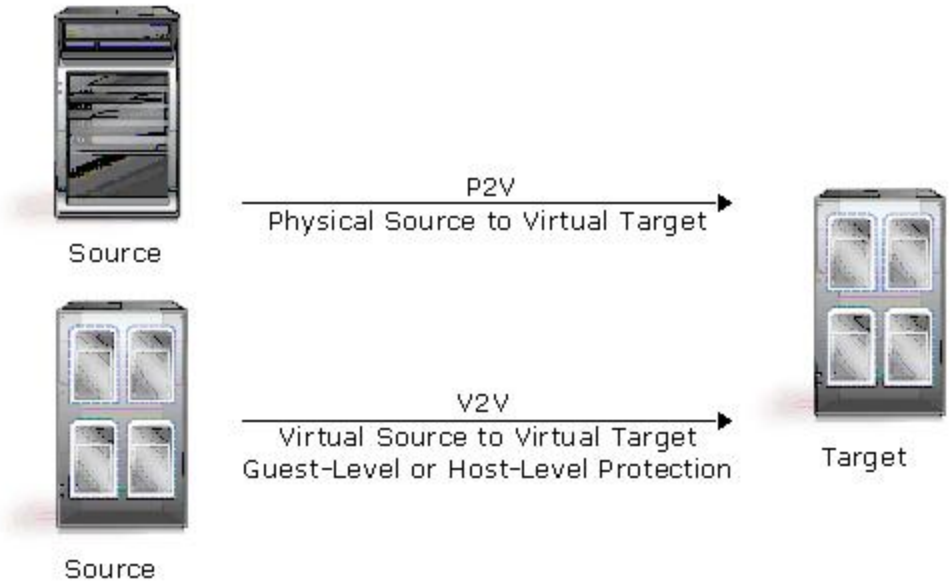

Cluster workloads

In a standard cluster configuration, a single copy of data resides on a SCSI disk that is shared between cluster nodes. Data is available without users knowing which node owns a cluster resource. MSCS handles failover between nodes of the cluster. By adding Double-Take Availability to this cluster environment, you can further protect your data by replicating the cluster data to a target. In the event the cluster fails, your cluster data will be available on the target.

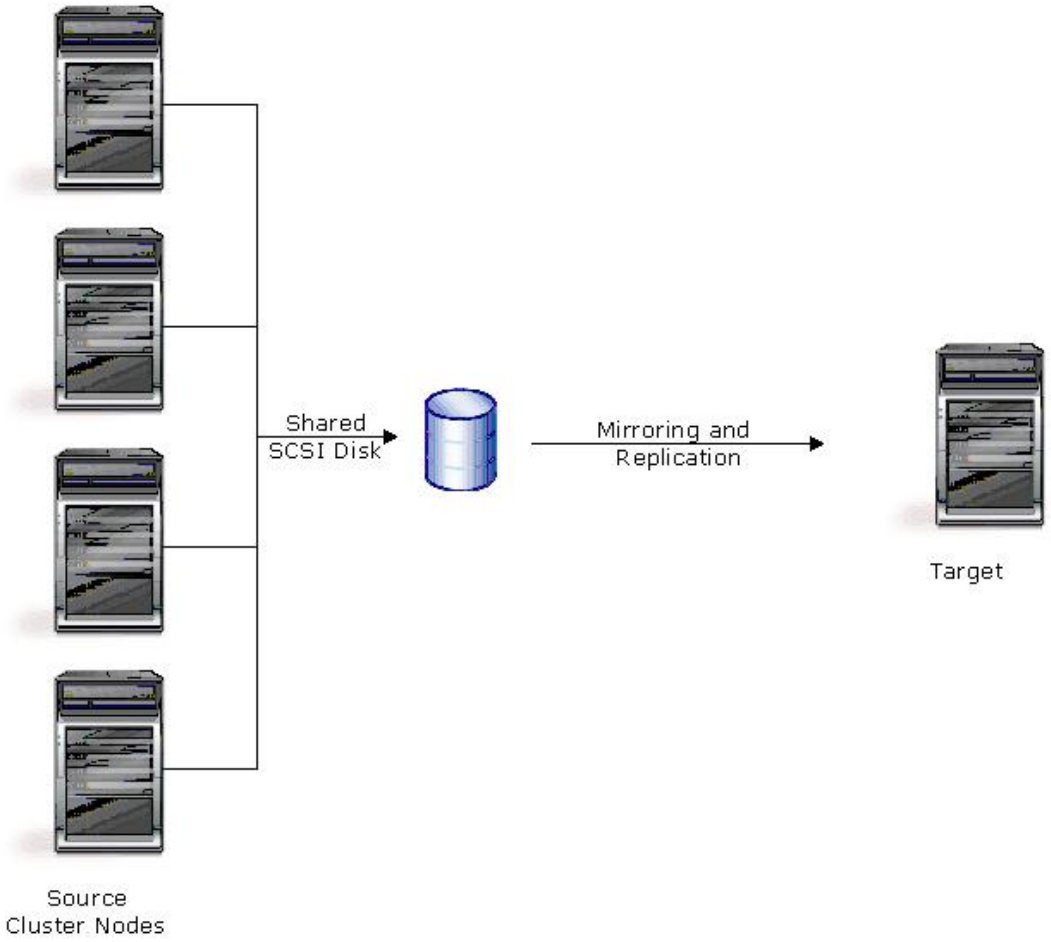

In a GeoCluster configuration, data is stored on volumes local to each node and replicated to each node in the cluster using Double-Take Availability. Resources and groups are handled in the same manner as a standard cluster. Instead of assigning one group by SCSI drive, you assign one group per logical volume. If a server, disk, group, or network interface should fail, MSCS relocates the failed group to another node, which contains the replicated copy of the data, thus maintaining availability.

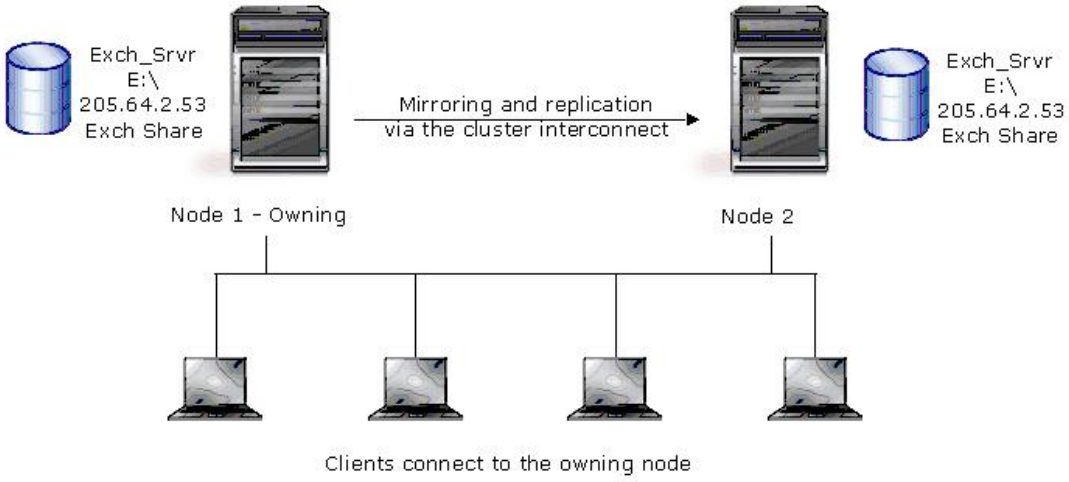

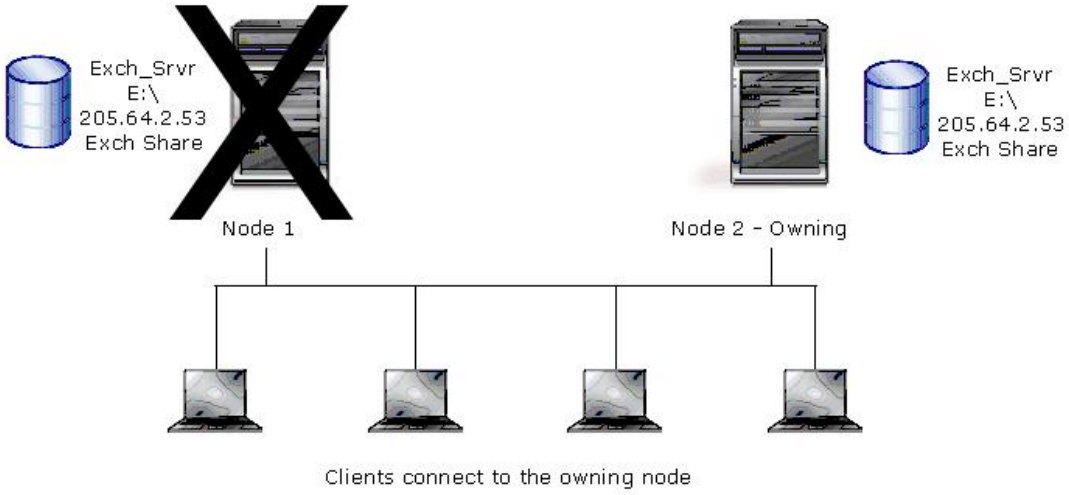

Supported configurations

Double-Take Availability is an exceptionally flexible product that can be used in a wide variety of network configurations. To implement Double-Take Availability effectively, it is important to understand the possible configuration options and their relative benefits. Double-Take Availability configuration options can be used independently or in varying combinations.

- One-to-one, active/standby
- One-to-one, active/active
- Many-to-one
- One-to-many
- Chained

One-to-one, active/standby

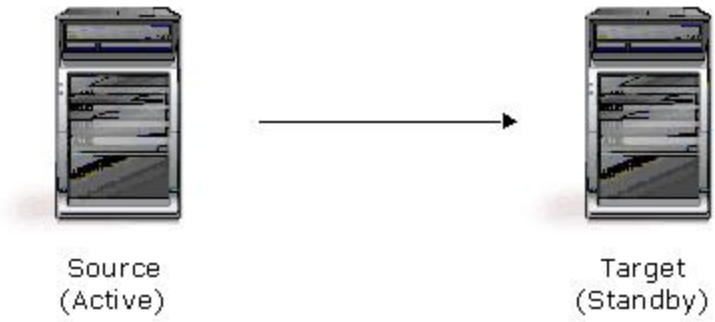

Description

One target server, having no production activity, is dedicated to support one source server. The source is the only server actively replicating data.

Applications

- This configuration is appropriate for offsite disaster recovery, failover, and critical data backup. This is especially appropriate for critical application servers such as Exchange, SQL Server, and web servers.
- This is the easiest configuration to implement, support, and maintain.

Considerations

- This configuration requires the highest hardware cost because a target server is required for every source server.
- You must pause the target when backing up database files on the target.

One-to-one, active/active

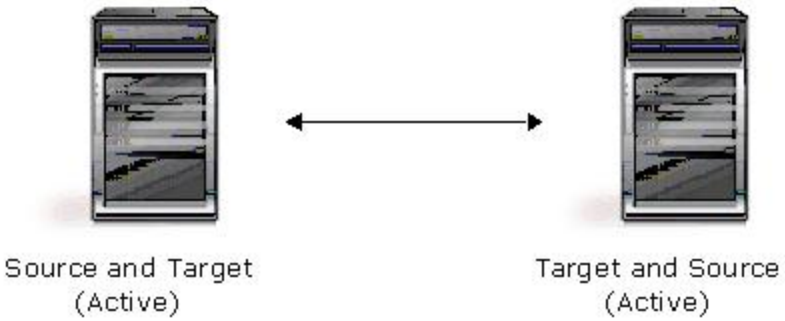

Description

Each server acts as both a source and target actively replicating data to each other

Applications

This configuration is appropriate for failover and critical data backup. This configuration is more cost-effective than the Active/Standby configuration because there is no need to buy a dedicated target server for each source. In this case, both servers can do full-time production work.

Considerations

- Coordination of the configuration of Double-Take Availability and other applications can be more complex than the one-to-one active/standby configuration.
- During replication, each server must continue to process its normal workload.
- Administrators must avoid selecting a target destination path that is included in the source's replication set. Any overlap will cause an infinite loop.
- To support the production activities of both servers during failover without reducing performance, each server should have sufficient disk space and processing resources.
- Failover and failback scripts must be implemented to avoid conflict with the existing production applications.
- You must pause the server when backing up database files.

Many-to-one

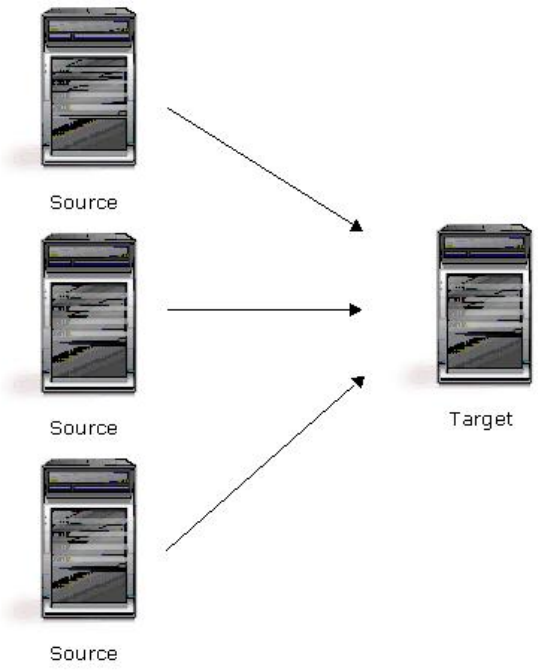

Description

Many source servers are protected by one target server.

Applications

This configuration is appropriate for offsite disaster recovery. This is also an excellent choice for providing centralized tape backup because it spreads the cost of one target server among many source servers.

Considerations

- The target server must be carefully managed. It must have enough disk space and RAM to support replication from all of the source systems. The target must be able to accommodate traffic from all of the servers simultaneously.
- If using failover, scripts must be coordinated to ensure that, in the event that the target server stands in for a failed server, applications will not conflict.
- You must pause the target when backing up database files on the target.

One-to-many

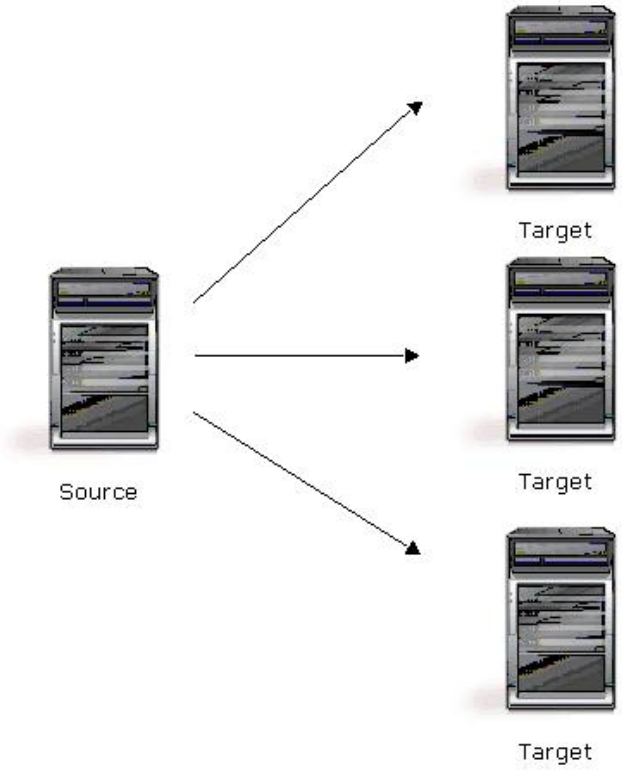

Description

One source server sends data to multiple target servers. The target servers may or may not be accessible by one another.

Applications

This configuration provides offsite disaster recovery, redundant backups, and data distribution. For example, this configuration can replicate all data to a local target server and separately replicate a subset of the mission-critical data to an offsite disaster recovery server.

Considerations

- Updates are transmitted multiple times across the network. If one of the target servers is on a WAN, the source server is burdened with WAN communications.
- You must pause the target when backing up database files on the target.

Chained

Description

The source servers sends replicated data to a target server, which acts as a source server and sends data to a final target server, which is often offsite.

Applications

This is a convenient approach for integrating local high availability with offsite disaster recovery. This configuration moves the processing burden of WAN communications from the source server to the target/source server. After failover in a one-to-one, many-to-one, or one-to-many configuration, the data on the target is no longer protected. This configuration allows failover from the first source to the middle machine, with the third machine still protecting the data.

Considerations

- The target/source server could become a single point of failure for offsite data protection.
- You must pause the target when backing up database files on the target.

Double-Take Availability requirements

Each Double-Take Availability server must meet minimum requirements. If you are protecting certain workloads, your servers may need to meet additional requirements. Verify that each server meets the general requirements and any additional requirements for your workload type. Additionally, the machine where you will be running the console must also meet several requirements.

- General source and target server requirements
- Full-server workload requirements
- Application workload requirements
- Virtual workload requirements
- Cluster workload requirements
- Double-Take Console requirements

General source and target server requirements

Verify that each server meets the following general source and target requirements, and then confirm your servers meet any additional requirements for your workload type.

- **Operating system**—There are different Double-Take Availability editions depending on the operating system you are using. Be sure you have the correct Double-Take Availability edition for your operating system.

  Note: Microsoft Server Core is supported for data workloads. See the specific requirements for the workload type you are protecting for additional Server Core requirements.

  Each of the Windows 2003 operating systems require Service Pack 1 or later.

  - Foundation Edition
  - Standard Edition
  - Advanced Edition
  - Premium Edition
  - Virtual Guest 5-Pack Edition
  - Virtual Host Standard Edition
  - Virtual Host Advanced Edition
  - Virtual Host Premium Edition

- **File system**—Double-Take Availability supports the same file system format that Microsoft supports: FAT, FAT32, and NTFS.
- **System memory**—There are different memory requirements depending on the operating system you are using. Be sure you have at least the minimum amount of memory for your environment. You may want to consider having at least the recommended amount of system memory.
  - **i386 operating systems**—The minimum system memory is 128 MB. The recommended system memory is at least 512 MB.
  - **x64 operating systems**—The minimum system memory is 512 MB. The recommended system memory is at least 1024 MB.
- **Disk usage**—The amount of disk space required for the Double-Take Availability program files is approximately 70 MB. You will need to verify that you have additional disk space for Double-Take Availability queuing, logging, and so on.

Additionally, on a target server, you need sufficient disk space to store the replicated data from all connected sources, allowing additional space for growth.

Note: The program files can be installed to any volume while the Microsoft Windows Installer files are automatically installed to the operating system boot volume.

- **Server name**—Double-Take Availability includes Unicode file system support, but your server name must still be in ASCII format. If you have the need to use a server's fully-qualified domain name, your server cannot start with a numeric character because that will be interpreted as an IP address.
- **Microsoft .NET Framework**—Double-Take Availability requires the Microsoft .NET Framework version 3.5 Service Pack 1. You can install this version from the Double-Take Availability CD, via a web connection during the Double-Take Availability installation, or from a copy you have obtained manually from the Microsoft web site.
- **Protocols and networking**—Your servers must meet the following protocol and networking requirements.
  - Your servers must have TCP/IP with static IP addressing.
  - By default, Double-Take Availability is configured for IPv6 and IPv4 environments, but the Double-Take service will automatically check the server at service startup and modify the appropriate setting if the server is only configured for IPv4.
  - IPv6 is only supported for Windows 2008 servers.
  - If you are using IPv6 on your servers, your clients must be run from an IPv6 capable machine.
  - In order to properly resolve IPv6 addresses to a hostname, a reverse lookup entry should be made in DNS.
- **Windows Management Instrumentation (WMI)**—Double-Take Availability is dependent on the WMI service. If you do not use this service in your environment, contact technical support.
- **Windows firewall**—The installation program will automatically attempt to configure ports 6320, 6325, 6330, and 6332 for Double-Take Availability. If you cancel this step, you will have to configure the ports manually.
- **Snapshots**—Double-Take Availability uses the Microsoft Volume Shadow Copy service for snapshot capabilities. To use this functionality, your servers must meet the following requirements.

- **Snapshot operating system**—Your servers must be running, at a minimum, Windows 2003 Service Pack 1. You should upgrade to Service Pack 2 or later so that several Microsoft patches that address memory leaks in the Volume Shadow Copy service are applied. If you do not have Service Pack 2 installed, you will need to review the patches available on the Microsoft web site and install those that correct the Volume Shadow Copy service memory leaks.
- **Snapshot file system**—Your servers must be using the NTFS file system. If you are using a FAT file system, the FAT volumes will not be included in the snapshot set, and when the snapshots are reverted, the FAT volume will not be time-consistent with the NTFS volumes.
- **Snapshot configuration**—If you have related data on different drives on your source (for example, an Exchange database on one drive and related log files on another), snapshots of each drive must be taken simultaneously so that the snapshots for each drive represent the same point in time. However, based on snapshot technology, different snapshots cannot be taken at the same time. To work around this limitation and guarantee data integrity on the target, you need to create a mount point, thus ensuring that one point-in-time consistent snapshot will be taken of both volumes at once. To create a mount point, create an empty folder on one of the drives. Using the **Disk Management** tool in the **Windows Computer Management** applet, remove the drive letter from the other drive. Then create a mount point by selecting **Mount in the following empty NTFS folder** and specifying the folder you just created. If you have multiple drives, create an empty folder on the drive for each of the other drives and create mount points to each of the folders. When the snapshot is taken of the drive, each mount point will be included in the snapshot. You will need to modify your applications to specify the new location for the files that are now on the mount point. For additional details on creating and using mount points, see your Windows reference manual. For details on modifying your applications, see your application reference manual.

Foundation Edition

Windows 2003 and 2003 R2 operating systems

- Storage Server Edition
- Small Business Server

Windows 2008 and 2008 R2 operating systems

- Storage Server Edition i386 and x64
- Small Business Server Standard and Premium x64
- Foundation Server with SAK
- Essential Business Server x64

Virtual server protection

There is no virtual server (guest-level or host-level) protection with the Foundation Edition.

Notes

You can install the Foundation Edition on a server running other server Windows operating systems, so you do not have to pay extra for an upgraded Double-Take Availability license. When the Foundation Edition is installed on a server running these other server Windows operating systems, the following limitations will apply. 1) The server will function in a target role only. 2) The target-only server can only protect a source that is running the Double-Take Availability Foundation Edition and one of the operating system editions listed for the Foundation Edition. 3) Full-server failover is the only supported method of failover.

Standard Edition

Windows 2003 and 2003 R2 operating systems

- Storage Server Edition
- Small Business Server
- Web Edition i386 and x64
- Standard Edition i386 or x64

Windows 2008 and 2008 R2 operating systems

- Storage Server Edition i386 and x64
- Small Business Server Standard and Premium x64
- Foundation Server
- Essential Business Server x64
- Web Server i386 and x64
- Standard Edition i386 and x64

Virtual server guest-level protection

The Standard Edition can run inside one virtual server to provide workload protection for data and/or applications running on the guest operating system.

Virtual server host-level protection protection

There is no virtual server host-level protection with the Standard Edition.

Advanced Edition

Windows 2003 and 2003 R2 operating systems

- Storage Server Edition
- Small Business Server
- Web Edition i386 and x64
- Standard Edition i386 or x64
- Enterprise Edition i386 and x64

Windows 2008 and 2008 R2 operating systems

- Storage Server Edition i386 and x64
- Small Business Server Standard and Premium x64
- Foundation Server
- Essential Business Server x64
- Web Server i386 and x64
- Standard Edition i386 and x64
- Enterprise Edition i386 and x64

Virtual server guest-level protection

The Advanced Edition can run inside one virtual server to provide workload protection for data and/or applications running on the guest operating system.

Virtual server host-level protection protection

There is no virtual server host-level protection with the Advanced Edition.

Notes

If you are using Hyper-V R2 with Cluster Shared Volumes (CSV), you can use Double-Take Availability within the guest to protect your workloads. However, you will be unable to protect the virtual machines at the host level using Double-Take Availability.

Premium Edition

Windows 2003 and 2003 R2 operating systems

- Storage Server Edition
- Small Business Server
- Web Edition i386 and x64
- Standard Edition i386 or x64
- Enterprise Edition i386 and x64
- Enterprise Itanium IA64
- Datacenter i386, x64, IA64

Windows 2008 and 2008 R2 operating systems

- Storage Server Edition i386 and x64
- Small Business Server Standard and Premium x64
- Foundation Server
- Essential Business Server x64
- Web Server i386 and x64
- Standard Edition i386 and x64
- Enterprise Edition i386 and x64
- Itanium Edition
- Datacenter Edition i386 and x64

Virtual server guest-level protection

The Premium Edition can run inside an unlimited number of virtual servers to provide workload protection for data and/or applications running on the guest operating system. These licenses can be spread across multiple hosts.

Virtual server host-level protection protection

There is no virtual server host-level protection with the Premium Edition.

Notes

If you are using Hyper-V R2 with Cluster Shared Volumes (CSV), you can use Double-Take Availability within the guest to protect your workloads. However, you will be unable to protect the virtual machines at the host level using Double-Take Availability.

Virtual Guest 5-Pack Edition

Windows 2003 and 2003 R2 operating systems

- Storage Server Edition
- Small Business Server
- Web Edition i386 and x64
- Standard Edition i386 or x64
- Enterprise Edition i386 and x64
- Enterprise Itanium IA64
- Datacenter i386, x64, IA64

Windows 2008 and 2008 R2 operating systems

- Storage Server Edition i386 and x64
- Small Business Server Standard and Premium x64
- Foundation Server
- Essential Business Server x64
- Web Server i386 and x64
- Standard Edition i386 and x64
- Enterprise Edition i386 and x64
- Itanium Edition
- Datacenter Edition i386 and x64

Virtual server guest-level protection

The Virtual Guest 5-Pack Edition can run inside five virtual servers to provide workload protection for data and/or applications running on the guest operating system. These licenses can be spread across multiple hosts.

Virtual server host-level protection protection

There is no virtual server host-level protection with the Virtual Guest 5-Pack Edition.

Notes

If you are using Hyper-V R2 with Cluster Shared Volumes (CSV), you can use Double-Take Availability within the guest to protect your workloads. However, you will be unable to protect the virtual machines at the host level using Double-Take Availability.

Virtual Host Standard Edition

Windows 2008 and 2008 R2 operating systems

- Standard Edition

Virtual server guest-level protection

There is no virtual server guest-level protection with the Virtual Host Standard Edition.

Virtual server host-level protection protection

The Virtual Host Standard Edition runs outside the virtual server to protect the hard disk (.vmdk or .vhd) and its associated files stored on the host operating system. This edition can protect five ESX hosts and an unlimited number of Hyper-V hosts.

Virtual Host Advanced Edition

Windows 2008 and 2008 R2 operating systems

- Standard Edition
- Enterprise Edition

Virtual server guest-level protection

There is no virtual server guest-level protection with the Virtual Host Advanced Edition.

Virtual server host-level protection protection

The Virtual Host Advanced Edition runs outside the virtual server to protect the hard disk (.vmdk or .vhd) and its associated files stored on the host operating system. This edition can protect ten ESX hosts and an unlimited number of Hyper-V hosts.

Virtual Host Premium Edition

Windows 2008 and 2008 R2 operating systems

- Standard Edition
- Enterprise Edition
- Datacenter Edition

Virtual server guest-level protection

There is no virtual server guest-level protection with the Virtual Host Premium Edition.

Virtual server host-level protection protection

The Virtual Host Premium Edition runs outside the virtual server to protect the hard disk (.vmdk or .vhd) and its associated files stored on the host operating system. This edition can protect an unlimited number of ESX hosts and an unlimited number of Hyper-V hosts.

Full-server workload requirements

If you will be protecting an entire server, the general source and target server requirements apply. However, keep in mind that a target server may meet these requirements but may not be suitable to stand-in for a source in the event of a source failure. See Finding a compatible target for additional information regarding an appropriate target server for your particular source.

Note: If you are using Small Business Server, you will need to check with your NAS vendor to verify if there are technical or license restrictions on failing over an image of a server to different hardware.

Microsoft Server Core is only supported in a Server Core to Server Core configuration.

There is one limitation for full-server protection of a cluster. Microsoft clusters identify disks by their disk signature, which is stored on the physical disk in the master boot record. When Double-Take Availabilityis used to failover a cluster node, the disk signature of the target will be different than the source and the cluster will fail to start. Therefore, Double-Take Availability does not natively support the full-server protection of Microsoft cluster nodes. However, you can alter the disk signature on the target manually. See the Microsoft Knowledge Base article 280425 for details on how to change the disk signature.

Application workload requirements

If you will be protecting an application, the general source and target server requirements apply. In addition, you must meet the application requirements below.

- **Server and network configuration**—Application protection requires the following server and network requirements.
  - The application program files must be installed in the same location on the source and target.
  - The drive letter(s) where the applications stores its data on the source must be the same on the target.
  - Single-label DNS domain names (those without a suffix such as .com, .corp, .net) are not supported.
  - In environments where the FIPS security policy is enabled, you must use impersonation, which requires the following.
  - The user running Double-Take Availability must have all appropriate rights to update the domain (that is, only impersonation is supported).
  - You must manually verify DNS rights by running the DFO utility with the /test parameter.
  - Microsoft Server Core is only supported for file server protection.
- **Verification**—If you want to use the Target Data Verification feature to confirm the integrity of your Exchange or SQL data on the target, you will need to install the Volume Shadow Copy Service SDK in the \windows\system32 directory on the target. You can download the SDK from the Microsoft website. Additionally, your Exchange version must be Exchange 2003 with service pack 1 or later.
- **Application Manager Console**—The following requirements and limitations apply to the Application Manager Console.
  - Ideally, you should run the console from a client machine or from the target.
  - Do not run the Application Manager Console from a domain controller.
  - If you are using a cluster configuration and run the Application Manager Console from a workstation operating system, you must have the Windows Administration Pack installed in order to have the cluster components installed on the workstation operating system.
- **Applications**—In addition to these general application requirements, you must also meet the requirements for the application you are protecting. See Exchange, SQL, SharePoint, BlackBerry, or File Server for those requirements.

Exchange protection requirements

In addition to the general application requirements, you must also meet the following requirements to protect Exchange.

- **Exchange versions**—Double-Take Availability can protect Microsoft Exchange 2003 or Exchange 2007, with the following requirements and limitations.
  - The version of Exchange on the source and target must be identical.
  - Double-Take Availability does not check the edition of Exchange 2007 (Enterprise or Standard). However, it is able to differentiate between service pack levels. If you have Exchange 2007 Enterprise on the source and Exchange 2007 Standard on the target, you will be limited to only failing over the number of databases or storage groups supported by Exchange 2007 Standard. See the Exchange Server 2007 Editions and Client Access Licenses information on the Microsoft website.
  - For Exchange 2007, in a consolidated role environment only the mailbox role is protected. The Hub Transport and Client Access roles are not protected or failed over because they are already installed on the target.
  - For Exchange 2007, replication and failover between a distributed role source configuration to a consolidated role target configuration is permitted as long as the source Mailbox Server role is installed on a standalone server or cluster with the other roles residing on different servers, and the target configuration is a standalone server with the Mailbox, Hub Transport, and Client Access roles installed. In these configurations, Double-Take Availability will not replicate any data associated with the Hub Transport/Client Access data, however, the target Hub Transport/Client Access roles function properly when failing over the source Mailbox role, allowing necessary operations to resume.
- **Server and network configuration**—The following requirements and limitations apply to your Exchange server and network configuration.
  - You can use a one-to-one configuration for Exchange protection. You cannot use a one-to-many, many-to-one, or chained configuration.
  - The source and target servers must be in the same root forest domain.
  - In a parent/child domain, at least one domain controller in the child domain must be designated as a global catalog server.
  - The target server cannot be a domain controller.
  - Exchange and a domain controller cannot be on the same node of a cluster.
  - Exchange 2003 on a domain controller is not a recommended configuration. However, if you must run Exchange 2003 on a domain controller, review

Microsoft Knowledge Base articles 822179, 332097, 305065, 304403, and 875427.

- The source and target servers must be part of the same Exchange Administrative Group.
- The Exchange configurations on the source and target servers must be identical for the following components: storage groups, location of storage groups (log and data files), log file prefixes, database locations (log and data files), Message Transfter Agent (MTA) location, and queue paths. If you are using like-named clusters, this requirement does not apply.
- Public folder replication issues may occur in environments where public folder replicas are shared between the source, target, and other Exchange server(s). In these environments, event IDs 3085 and 3092 are logged to the Application event log after failback. If you are using Exchange 2003, no additional steps are required. If you are using Exchange 2007, you must run the Application Manager Console on a machine, preferable the target server, that has PowerShell and the Exchange Management snap-in installed in order to address the issue.
- Before you attempt to protect your Exchange application, you may want to complete the following tasks to verify that the environment is properly set up.
  - With both Exchange servers online, use **Active Directory Users and Computers** to move an existing user from the source to the target and then back to the original source.
  - Verify that you can create a new user on the target.
  - To verify connectivity, create an Outlook profile for the new user on a client machine and connect to the target.
  - If /domainprep has not been run in an Exchange 2007 environment, users will not be failed over and SPNs will not be updated during failover due to access denied errors. To fix this issue, run setup with the /domainprep parameter in the environment.

- **Cluster protection**—Cluster to cluster and cluster to standalone configurations are supported. A standalone to cluster configuration is not supported. In addition, the following limitations apply.
  - If you are using Exchange 2007, only the mailbox role is protected.
  - Exchange and the domain controller cannot be on the same node in the cluster.
  - Exchange must be installed in a unique group, not in the cluster group.
  - If you are using a Windows Server 2008 cluster, the Application Manager must be running a machine with the Windows 2008 failover cluster management tools.

- **Like-named cluster protection**—If you are using Exchange 2003, you can protect a cluster with a like-named cluster, also known as a standby cluster. Double-Take Availability will move the Exchange virtual server from the source cluster to the target cluster. The process of moving users and public folders from one server to another is not needed because users will continue to use the same mail store on the target as they were on the source. The like-named cluster environment must meet the following requirements.
  - All nodes on the source and target clusters must have the same Exchange version and service pack.
  - The source and target clusters must have the same Exchange resource group name.
  - The target resource group only needs to contain physical disk resources, however they must use the same drive letters that are used by the physical disk resources on the source.
  - Application Manager will create temporary name and IP address resources on the target cluster for Double-Take Availability mirroring and replication. A DNS entry is created based on the target's owning node DNS server. If the source and target owning nodes are configured to use different DNS servers, this can cause issues when enabling protection. If you have issues enabling protection in this configuration, verify that the source's owning node DNS server is correctly set up to receive DNS zone updates from the target's owning node DNS server, or reload the forward and reverse zones using the dnsmgmt utility.
  - The user configuring the like-named cluster protection in the Application Manager must be a member of the local administrator and Double-Take Admin groups on all cluster nodes. Additionally, the user must be delegated Full Exchange Administrator access through Exchange System Manager. Optionally, the user must be a member of the DnsAdmins group if you want to be able to update the Time to Live DNS attribute.
  - If you have multiple Exchange virtual servers, you can configure multiple like-named cluster protection connections, or you can failover multiple Exchange virtual servers to pre-existing Exchange virtual servers on the target.
- **Security**—By default, the Double-Take service is configured to use the local system account. If you are protecting Exchange, you cannot change this configuration.

Note: If you are protecting Exchange in a 2008 R2 domain where the domain functional level is set to R2, you must grant two levels of access to the local system account on the target. First, grant the **Administer information store**

control to the target in the **Configuration Naming Context** in order to move users during failover. Second, grant **Full control** to the target in the **Domain Naming Context** in order to move Service Principal Names, which will allow users to access their e-mail after a failover.

- **Application Manager Console**—The following limitations apply to the Application Manager Console when protecting Exchange.
  - If you are using Exchange 2007, the console may be run on a workstation provided that the Exchange 2007 Management Tools are installed prior to installing Double-Take Availability.
  - The machine running the console must have access to the domain in which the Exchange servers are located.

SQL protection requirements

In addition to the general application requirements, you must also meet the following requirements to protect SQL.

- **SQL versions**—Double-Take Availability can protect Microsoft SQL Server or Express 2000 with Service Pack 4 or later, Server or Express 2005, or Server or Express 2008, with the following requirements and limitations.
  - If you are using Windows 2008, you can protect SQL Server or Express 2005 or 2008. SQL Server or Express 2000 is not supported on Windows 2008.
  - You should use the same version, service pack, and architecture (32-bit or 64-bit) of SQL Server on both the source and target servers. The only exceptions is in database only protection mode you. In this case, you can use a newer version of SQL Server on the target server, or you may have a 32-bit source and a 64-bit target. For example, you may want migrate from SQL Server 2000 on the source to SQL Server 2005 on the target, or migrate data from a 32-bit source to a 64-bit target. However, you cannot failback using different versions of SQL Server on the source and target.
  - If you are using SQL Express 2000 or Express 2005, named pipes and TCP/IP need to be added to the enabled protocols to accept remote connections. By default, these are disabled.
  - For Express 2000, you must run svrnetcn.exe, which is located in the C:\Program Files\Microsoft SQL Server\80\Tools\Binn directory.
  - For Express 2005, you must launch the SQL Server Configuration Manager. Expand **SQL Server 2005 Network Configuration**, and under **Protocols for MSSQLSERVER** enable **Named Pipes** and **TCP/IP**.
- **Server and network configuration**—The following requirements and limitations apply to your SQL server and network configuration.
  - You can use a one-to-one configuration for SQL protection. You can also use a many-to-one configuration, however protection will be database mode only. You cannot use a one-to-many or chained configuration.
  - The source and target servers should be in the same domain. If they are not, the SQL Server service on both the source and target servers must be configured to start with the same domain user account.
  - If your source and target are in a workgroup, make sure the source server's NIC does not register the IP addresses with DNS.
  - In order to protect SQL named instances, both the source and target servers must have named instances with the same name installed prior to configuring protection.

- If you are using SQL 2005 and are using a domain service account that is not in the domain or local Admins security group, the replicated databases will not mount on the target because of security restrictions at the file system level. You need to place the SQL 2005 service account in the local Admins group on the target.
- If you are using SQL 2008, local group permissions are not copied to the target during failover. You will need to use a domain Admin account as the SQL service account or manually copy the group permissions to the target.
- Double-Take Availability does not support a SQL default instance that is using non-default ports.
- Cluster to cluster and cluster to standalone configurations are supported for SQL Server 2005 and 2008 only.
- If you are protecting a cluster configuration, ideally you should have only one instance of SQL Server per owning node on your source and target cluster. This decreases the risk of problems when attempting to re-enable protection after failover and failback. The following problems may occur if multiple instances reside on the same owning node.
- After failover and failback, you will be able to re-enable protection on the first instance, but when you try to protect subsequent instances, after selecting the source instance, Application Manager will erroneously show that the instance is already protected. Contact technical support to obtain instructions for fixing this issue.
- If one instance is failed over, replication will stop on the other protected instances. On the other instances that are not failed over, you will need to perform a difference mirror in order to resume protection.

- **Application Manager Console**—The following limitations apply to the Application Manager Console when protecting SQL.
  - If you are using a cluster configuration, you should run the console from a target node.
  - If you want to run the console from a Vista client to protect a cluster, you will need to install the Microsoft Remote Server Administration Tools (RSAT). Installing this package will allow you to install the Failover Cluster Manager component on a Vista client so it can communicate with and administer Windows 2008 clustered environments.

SharePoint protection requirements

In addition to the general application requirements, you must also meet the following requirements to protect SharePoint.

- **SharePoint versions**—Double-Take Availability can protect Windows SharePoint Services (WSS) version 3 or Microsoft Office SharePoint Server (MOSS) 2007. If you are using Windows 2008, you must have Service Pack 1 for each of these SharePoint versions.
  - Application support for SharePoint is for SQL instance mode protection. Database-only protection mode is not available.
  - Only target web servers running a version of SharePoint that is identical to what is installed on the source web front-end can be extended into the source SharePoint configuration.
  - The SharePoint Admin account used to install SharePoint on the source web front-end is required to extend a target web server into the SharePoint configuration.
- **Server and network configuration**—The following requirements and limitations apply to your SharePoint server and network configuration.
  - You can use a one-to-one configuration for SharePoint protection. You cannot use a one-to-many, many-to-one, or chained configuration.
  - Each source and target must have SQL Server 2000, 2005 or 2008 installed.
  - You will need to open port 6350 on each source and target for SharePoint communication
  - SharePoint protection is currently only supported in a flat domain structure.

BlackBerry protection requirements

In addition to the general application requirements, you must also meet the following requirements to protect BlackBerry.

- **BlackBerry versions**—Double-Take Availability can protect BlackBerry Enterprise Server 4.1.4 through 4.1.6 for Microsoft Exchange.
- **Server and network configuration**—The following requirements and limitations apply to your BlackBerry server and network configuration.
  - Only the Windows 2003 operating systems listed in the general source and target server requirements can be used to protect BlackBerry.
  - Each source and target must have Microsoft Exchange 2003 or 2007.
  - The version of Exchange on the source and target must be identical.
  - Each source and target must have SQL Server 2000 with Service Pack 4, 2005, or 2008 installed.
  - Double-Take Availability does not support database only mode. Only the instance of SQL where BlackBerry is installed will be protected.
  - You should use the same version, service pack, and architecture (32-bit or 64-bit) of SQL Server on both the source and target servers. The only exceptions is in database only protection mode you. In this case, you can use a newer version of SQL Server on the target server, or you may have a 32-bit source and a 64-bit target. For example, you may want migrate from SQL Server 2000 on the source to SQL Server 2005 on the target, or migrate data from a 32-bit source to a 64-bit target. However, you cannot failback using different versions of SQL Server on the source and target.

File Server protection requirements

In addition to the general application requirements, you must also meet the following requirements to protect a file server.

- You can use a one-to-one configuration for file server protection. You cannot use a one-to-many, many-to-one, or chained configuration.
- The target must be a dedicated, stand-by server which does not host any critical applications.
- During failback, the Server service is restarted, which may also restart any dependent services.
- File server protection is currently only supported in a flat domain structure.
- Microsoft Server Core is supported for file server protection.

Virtual workload requirements

When you are protecting virtual server workloads, the general source and target server requirements still apply, however, the requirements and limitations are more strict depending on your virtual configuration. Select a link to see the requirements that correspond with your virtual configuration.

- **Physical or virtual (guest level) to Hyper-V virtual**—If your source is a physical or virtual server, and you want to protect the volumes from the physical server or the volumes from within the virtual guest operating system, and your target is a virtual server on a Hyper- V server, then see the Physical or virtual to Hyper-V requirements.
- **Physical or virtual (guest level) to ESX virtual**—If your source is a physical or virtual server, and you want to protect the volumes from the physical server or the volumes from within the virtual guest operating system, and your target is a virtual server on an ESX server, then see the Physical or virtual to ESX requirements.
- **Hyper-V virtual (host level) to Hyper-V virtual**—If your source is a virtual server on a Hyper-V server, and you want to protect the host-level virtual disk files (.vhd files), and your target is a virtual server on a Hyper-V server, then see the Hyper-V to Hyper-V requirements.
- **ESX virtual (host level) to ESX virtual**—If your source is a virtual server on an ESX server, and you want to protect the host-level virtual disk files (.vmdk files), and your target is a virtual server on an ESX server, then see the ESX to ESX requirements.

Physical or virtual to Hyper-V requirements

Use these requirements if your source is a physical or virtual server, you want to protect the volumes from the physical server or volumes from within the virtual guest operating system, and your target is an automatically provisioned virtual server on a Hyper-V server. This means Double-Take Availability will automatically create the virtual server on the Hyper-V target if failover is triggered.

- **Source server**—The source server can be any physical or virtual server running any of the operating systems listed in the General source and target server requirements. However, if you are using a Windows 2003 operating system, you must have Service Pack 2 which is required for Hyper-V Integration Services.
- **Target server**—The target server can be any Windows 2008 or 2008 R2 operating system from the supported server operating systems that has the Hyper-V role enabled. In addition, you can use Hyper-V Server 2008 R2 or Server Core 2008 R2 with the Hyper-V role enabled. (Hyper-V Server 2008 and Server Core 2008 are not supported.)

Physical or virtual to ESX requirements

Use these requirements if your source is a physical or virtual server, you want to protect the volumes from the physical server or volumes from within the virtual guest operating system, and your target is an automatically provisioned virtual server on an ESX server. This means Double-Take Availability will automatically create the virtual server on the ESX target if failover is triggered.

- **Source server**—The source server can be any physical or virtual server running any of the operating systems listed in the General source and target server requirements.
- **ESX server**—The ESX server that will host your target can be any of the following operating systems. Note that ESX is commonly referred to as the Classic edition and ESXi as the Embedded and Installable edition.
  - ESX 3.5.x or ESXi 3.5.x Standard, Advanced, or Enterprise
  - ESX 4.0 or ESXi 4.0 Standard, Advanced, Enterprise, or Enterprise Plus

  ---

  Note: If you are using the Standard edition of ESX 4.0 or ESXi 4.0, you must have update 1 or later.

  If your source is a Windows 2008 R2 server, your ESX server must have version 3.5 update 5 or later or ESX 4.0 update 1 or later.

- **VirtualCenter**—If you are using ESXi, VirtualCenter 2.5 or later is required. Although VirtualCenter is not required for the ESX versions, if you are using it, then you must use version 2.5 or later.

  Note: VMotion is only supported if you are using VirtualCenter.

- **Target server or virtual recovery appliance**—The target server must be an existing virtual machine, known as a virtual recovery appliance, running on the ESX server that meets the following requirements.
  - The virtual recovery appliance can be any of the operating systems listed in the General source and target server requirements.
  - The virtual recovery appliance must have the same or newer operating system than the source (not including service pack level).

    - The virtual recovery appliance must have Double-Take Availability installed and licensed on it.
    - When you establish protection, the virtual recovery appliance will create a new virtual server, mount disks, format disks, and so on. If failover occurs, the new virtual machine is detached from the virtual recovery appliance and powered on. Once the new virtual machine is online, it will have the identity, data, and system state of the source. Since the virtual recovery appliance maintains its own identity, it can be reused for additional failovers.
- **Server and network configuration**—The source and target cannot both be domain controllers. Only one or the other can be a domain controller.
    - Ideally, the target should not be a domain controller or host any functionality (file server, application server, and so on) because the functionality will be removed when failover occurs.
    - If your source is a domain controller, it will start in a non-authoritative restore mode after failover. This means that if the source was communicating with other domain controllers before failover, it will require one of those domain controllers to be reachable after failover so it can request updates. If this communication is not available, the domain controller will not function after failover. If the source is the only domain controller, this is not an issue.

Hyper-V to Hyper-V requirements

Use these requirements if your source is a virtual server on a Hyper-V server, you want to protect the host-level virtual disk files, and your target is a virtual server on a Hyper-V server.

- **Source and target operating system**—Your source and target servers can be any Windows 2008 or 2008 R2 operating system from the supported server operating systems that has the Hyper-V role enabled. In addition, you can use Hyper-V Server 2008 R2 or Server Core 2008 R2 with the Hyper-V role enabled. (Hyper-V Server 2008 and Server Core 2008 are not supported.)
- **Source and target configurations**—You can use one-to-one, many-to-one, or one-to-many configurations, however, you cannot use a chained configuration in a Hyper-V to Hyper-V to Hyper-V scenario.
- **Hyper-V configuration**—The following limitations apply to the virtual machines on the source and target Hyper-V servers.
  - The virtual machines must be in their own home folder that is not shared by any other virtual machines.
  - The virtual machines cannot be created in or replicated to the Hyper-V system default folder.
  - The virtual machines' snapshot folder must be unique to each virtual machine, they cannot be in the Hyper-V system default folder, and they cannot be changed once protection has been established.
  - The virtual machines cannot use raw, pass-through, or differencing disks.
- **Clusters**-Clustered Hyper-V hosts are not supported.
- **Ports**—In addition to the standard Double-Take Availability ports that must be opened, you must also open port 135 for communication between the client and the servers.

ESX to ESX requirements

Use these requirements if your source is a virtual server on an ESX server, you want to protect the host-level virtual disk files, and your target is a virtual server on an ESX server.

- **ESX servers**—The ESX servers that will host your source and target can be any of the following operating systems. Note that ESX versions listed are commonly referred to as the Classic edition.
  - ESX 3.5.x Standard, Advanced, or Enterprise
  - ESX 4.0.x Standard, Advanced, Enterprise, or Enterprise Plus

  ---

  Note: If your source and target are running on different versions of ESX, ideally, the target should be a newer version of ESX. If your source must be a newer version of ESX than the target, you must take into consideration ESX features that are supported on the newer version on the source (like VmxNet enhanced or additional virtual NICs available) that will not be supported on the earlier target version of ESX

  .

- **VirtualCenter**—Although VirtualCenter is not required, it is recommended. If you are using it, then you must use version 2.5 or later.

  Note: VMotion is only supported if you are using VirtualCenter.

  Do not use MSDE for the VirtualCenter database.

- **Double-Take Availability for VMware Infrastructure service and console**—The service and console can be run from a physical or virtual machine running Windows 7, Windows Vista, or Windows XP, or any of the supported server operating systems.
- **Ports**—In addition to the standard Double-Take Availability ports that must be opened, you must also open ports 22, 443, and 6331 on the source and target servers.
- **E-mail notification**—In order to use automatic e-mail notification, you must add VI_Service.exe to the exception list of most anti-virus and spam filters. In addition, you may need to open port 25 in your anti-trust software to allow SMTP e-mail.

Cluster workload requirements

- **Standard cluster**—If you will be protecting a standard cluster configuration, with a shared SCSI disk, there are two additional requirements.
  - **Network**—The cluster's private network should be a unique subnet so that Double-Take Availability will not attempt to use an unreachable private network.
  - **Volumes**—The source and target should have identical drive mappings.
- **GeoCluster**—If you will be protecting a GeoCluster configuration, where data is stored on volumes local to each node, there are several additional requirements.
  - **Operating system**—The operating systems are limited to the Windows Enterprise, Itanium, and Datacenter editions listed in the source and target server requirements.
  - **Volumes**—The source and target should have identical drive mappings.
  - **Hardware**—Intel-based hardware is required. Microsoft support for MSCS and MSCS-based Microsoft applications requires that the cluster configuration appear on the Microsoft Hardware Compatibility List under category Cluster.
  - **Cluster Network Name**—Double-Take Availability does not handle dynamic changes to the cluster network names, the names assigned to the routes for network traffic. If a network name is changed for a network that is used by a GeoCluster Replicated Disk resource, the resource must be taken offline, the resource's network property must be changed, and then the resource must be brought back online.
  - **Protocols and Networking**—All of the networking requirements below must be met.
    - TCP/IP connection between nodes on the same logical IP subnet
    - Your network can contain direct LAN connections or VLAN technology.
    - The maximum round trip latency between nodes should be no more than ½ second.
    - Multiple networks are recommended to isolate public and private traffic.

Double-Take Console requirements

There are various Double-Take Availability consoles, many of which are being phased out over time. To help consolidate the consoles and help you locate the necessary workflows to complete your work, use the console called Double-Take Console. You must meet the following requirements for the Double-Take Console.

- **Operating system**—The Double-Take Console can be run from a source or target. It can also be run from a 32-bit or 64-bit physical or virtual machine running Windows 7, Windows Vista, or Windows XP Service Pack 2 or later.
- **Microsoft .NET**—The Microsoft .NET Framework version 3.5 Service Pack 1 is required to run the console.
- **Screen resolution**—For best results, use a 1024x768 or higher screen resolution.

Installation

Before beginning the installation, review the Double-Take Availability requirements and Installation and upgrade notes. Use the installation instructions appropriate for the type of workload you are protecting

- **Data, full-server, and application workloads**—If you are protecting a data workload, full-server workload, or an application workload, you can install using the standard installation program or the command-line automatic installation process.
- **Virtual workloads**—You can also use the standard installation program or the command-line automatic installation process if you are protecting a virtual workload, however, if your source and target are both virtual machines on an ESX server and you want to protect the host-level virtual disk files, you need to install Double-Take Availability for VMware Infrastructure.
- **Clusters**—If you are protecting a cluster, the installation will depend on the type of cluster configuration you are using. For standard clusters with a shared SCSI disk, use the standard installation program or use the command-line automatic installation process. For a GeoCluster configuration, where data is stored on volumes local to each node and replicated to each node in the cluster, you will need to configure your cluster appropriately and install Double-Take Availability.

Installation and upgrade notes

Review the following installation and upgrade notes before beginning your installation or upgrade.

- Since Double-Take Availability 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 Windows reference manuals. Make sure that you select the option to backup the registry when building the repair disks.
- Because Double-Take Availability has operating system dependent files, if you are upgrading your operating system (to a new major version, not a service pack) and have Double-Take Availability installed, you must remove Double-Take Availability prior to the operating system upgrade. Uninstall Double-Take Availability, perform the operating system upgrade, and then reinstall Double-Take Availability.
- If you are installing to a drive other than the drive which contains your system TEMP directory, the Microsoft Windows Installer will still load 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 may need to change where it is located.
- During installation, a file called dtinfo.exe is installed to the Double-Take Availability installation directory. This program can be run to collect configuration data for use when reporting problems to technical support. It gathers Double-Take Availability log files; Double-Take Availabilityand system registry settings; network configuration information such as IP, WINS, and DNS addresses; and other data which may be necessary for customer support to troubleshoot issues. After running the executable, a zip file is automatically created with the information gathered.
- Double-Take Availability 5.2 is interoperable back to version 5.0 but is restricted to the following limitations. The Double-Take Availability clients can only control the same or older releases. To accommodate rolling upgrades, older sources can connect to newer targets, but newer sources cannot connect to older targets.
    - **5.0 client**—Supports 5.0 source and target, but does not support 5.1 or 5.2 source or target
    - **5.1 client**—Supports 5.0 or 5.1 source and target as long as the target is the same or newer than the source, but does not support 5.2 source or target
    - **5.2 client**—Supports 5.0, 5.1, or 5.2 source and target as long as the target is the same or newer than the source

- When performing a rolling upgrade, update the target servers first. After the upgrade is complete, the sources will automatically reconnect to the targets. Upgrade the sources when convenient.
- If you are using a chained configuration, update the last target first, then update the middle server acting as both a source and target, and update the production source last.
- If you are using a configuration where the source is an older version than the target, you will not be able to restore from the newer version target back to the older version source. You must upgrade the source to the same version as the target before restoring.
- Use the following procedure to upgrade Double-Take Availability on a cluster. If both your source and target are clusters, use the following procedure on the target cluster first, then on the source.
    1. Move all cluster resources to one node.
    2. Upgrade to the new version of Double-Take Availability on all of the other nodes.
    3. Move the cluster resources to one of the upgraded nodes.
    4. Upgrade to the new version of Double-Take Availability on the last node.
    5. If needed, move the cluster resources to the desired nodes.
- If you have protected clusters with an earlier version of the Application Manager, you should disable protection before upgrading to this version of Double-Take Availability.

Installing or upgrading Double-Take Availability

Use these instructions to install Double-Take Availability or upgrade an existing Double-Take Availability installation.

1. Close any open applications.
2. Start the installation program using the appropriate instructions, depending on your media source.
   - **CD**—Load the Double-Take Availability CD into the local CD-ROM drive. If auto-run is enabled, the installation program will start automatically. To manually start the program, select **Start**, **Run** and specify <cd_ drive>:\autorun.exe.
   - **Web download**—Launch the .exe file that you downloaded from the web.
3. When the installation program begins, the Double-Take Software Setup Launcher appears allowing you to install software and view documentation for various applications from Double-Take Software. The listed applications will depend on which products are included on the CD or in the web download. To install Double-Take Availability, select **Double-Take Availability** from the list of products. Under **Product Installs**, select Double-Take Availability.
4. Depending on your version of Windows and the components you have installed, you may see an initial screen indicating that you should install or enable Microsoft .NET Framework. If you do not see this screen, your server already has the appropriate version of Microsoft .NET. You should install or enable Microsoft .NET before installing Double-Take Availability. Click **Yes** to install Microsoft .NET. Click **No** to continue without installing .NET.
5. When the Double-Take Availability installation begins, you will be given the opportunity to check for a more recent version of the software.
   - If you do not want to check for a later version, select **No** and click **Next**.
   - If you want to check for a later version, select **Yes** and click **Next**. The installation program will establish an Internet connection from your server to the Double-Take Software web site.
   - If later versions are found, they will be listed. Highlight the version you want and either download that version and install it automatically or download that version and exit the installation. (If you exit the installation, you can run the updated installation later directly from the location where you saved it.)
   - If no later versions are found, continue with the current installation.
   - If an Internet connection cannot be established, continue with the current installation or install a previously downloaded version.

6. Review and accept the Double-Take Software license agreement to continue with the installation program. Click **Next** to continue.
7. Select the type of installation you would like to perform on this machine.
    - **Client and Server Components**—This option installs both the client and server components. The server components are required for systems that will function as a source or target. The server requires an activation code for the service to run. The client does not require an activation code, but it is required to administer this and other servers throughout the organization.
    - **Client Components Only**—This option installs only the client components. The client components do not require an activation code, but are required to administer servers throughout the organization.
    - **Server Components Only**—This option installs only the server components. The server components are required for systems that will function as a source or target. The server requires an activation code for the service to run.
8. If desired, specify where the Double-Take Availability files will be installed.
9. Click **Next** to continue.
10. You will be prompted to enter your activation code information. Your **Activation Code** is a 24-character, alpha-numeric activation code which applies the appropriate license to your installation. You must have a valid activation code to use Double-Take Availability. Enter your code and click **Add**.
11. Click **Next** to continue.
12. The next screen will depend on the activation code you entered.
    - If you have entered a valid activation code, you will be prompted to confirm the code. Click **Next** to continue the installation.
    - If you have entered an invalid activation code, you will be prompted that the code is incorrect and that the source and target modules will not load. Click **Back** and reenter your activation code.
    - If you have entered an evaluation activation code, the expiration date will be displayed and you will be prompted that the source and target modules will not load after that date. Click **Next** to continue the installation. You must update the activation code to a valid one before the expiration date, otherwise, on the expiration date, Double-Take Availability functionality will be disabled.
13. Double-Take Availability uses system memory to store data in queues. Specify the maximum amount of system memory to be used for the Double-Take Availability queues and click **Next** to continue.

    If you set the system memory queue lower, Double-Take Availability will use less system memory, but you will queue to disk sooner which may impact system

performance. If you set it higher, Double-Take Availability will maximize system performance by not queuing to disk as soon, but the system may have to swap the memory to disk if the system memory is not available. In general, the amount of memory Double-Take Availability and other applications on the server are configured to use should be less than the amount of physical memory on the system to prevent low memory conditions.

14. When the Double-Take Availability system memory queue is exhausted, Double-Take Availability will queue to disk. Specify the size and location of the disk queue. (See Queuing data for guidelines on selecting an appropriate location.) By default, the disk space is set to **Unlimited** which will allow the queue usage to automatically expand whenever the available disk space expands. Click **Next** to continue.
15. The Double-Take Availability security information screen appears next. Review this information and click **Next** to continue with the installation.
16. If you are satisfied with the selections you have made and are ready to begin copying the Double-Take Availability files, click **Install**.
17. During the installation, you may be prompted to add an exception to the Windows Firewall for Double-Take Availability. Click **OK** to add the port exception. If you **Cancel** the port modification, you will have to manually modify your firewall settings for Double-Take Availability processing.
18. After the files have completed copying, click **Finish** to exit the installation program.

Installing or upgrading Double-Take for VMware Infrastructure

In a typical Double-Take Availability installation, you install the server components on each source and target server. However, when you are protecting the host-level virtual disk files (the .vmdk files) from an ESX source to an ESX target, you need to install the server components on another machine. The Double-Take Availability for VMware Infrastructure service will run from this machine and will communicate with both the source and target servers to handle replication. You may want to consider protecting this machine with Double-Take Availability to keep the machine available because if it becomes unavailable, Double-Take Availability for VMware Infrastructure will be unable to replicate data between the source and target.

In addition to the Double-Take Availability for VMware Infrastructure service, you will need to install the Double-Take Availability for VMware Infrastructure console. The console can be installed on the same machine as the service or it can be installed on a different machine.

Use these instructions to install or upgrade an existing Double-Take Availability for VMware Infrastructure installation.

1. Close any open applications.
2. Start the installation program using the appropriate instructions, depending on your media source.
   - **CD**—Load the Double-Take Availability CD into the local CD-ROM drive. If auto-run is enabled, the installation program will start automatically. To manually start the program, select **Start**, **Run** and specify <cd_ drive>:\autorun.exe.
   - **Web download**—Launch the .exe file that you downloaded from the web.
3. When the installation program begins, the Double-Take Software Setup Launcher appears allowing you to install software and view documentation for various applications from Double-Take Software. The listed applications will depend on which products are included on the CD or in the web download. To install Double-Take Availability for VMware Infrastructure, select **Double-Take Availability** from the list of products. Under **Product Installs**, select **Double-Take Availability for VMware Infrastructure**.
4. Depending on your version of Windows and the components you have installed, you may see an initial screen indicating that you should install or enable Microsoft .NET Framework. If you do not see this screen, your server already has the appropriate version of Microsoft .NET. You should install or enable Microsoft .NET before installing Double-Take Availability for VMware Infrastructure. Click **Yes** to install Microsoft .NET. Click **No** to continue without installing .NET.

5. When the Double-Take Availability for VMware Infrastructure installation begins, click **Next**.
6. You will be given the opportunity to check for a more recent version of the software.
    - If you do not want to check for a later version, select **No** and click **Next**.
    - If you want to check for a later version, select **Yes** and click **Next**. The installation program will establish an Internet connection from your server to the Double-Take Software web site.
    - If later versions are found, they will be listed. Highlight the version you want and either download that version and install it automatically or download that version and exit the installation. (If you exit the installation, you can run the updated installation later directly from the location where you saved it.)
    - If no later versions are found, continue with the current installation.
    - If an Internet connection cannot be established, continue with the current installation or install a previously downloaded version.
7. Review and accept the Double-Take Software license agreement to continue with the installation program. Click **Next** to continue.
8. Specify where the Double-Take Availability for VMware Infrastructure files will be installed. Click **Next** to continue.
9. Select the type of installation you would like to perform on this machine.
    - **Client and Server**—This option installs both the Double-Take Availability for VMware Infrastructure service and console components. The service components are required to control replication between a source and target. The service requires an activation code. The console does not require an activation code, but it is required to administer the Double-Take Availability for VMware Infrastructure service.
    - **Client**—This option installs only the Double-Take Availability for VMware Infrastructure console components. The console does not require an activation code, but it is required to administer the Double-Take Availability for VMware Infrastructure service.
10. Click **Next** to continue.
11. You will be prompted to enter your user name, company name, and activation code information. Your **Activation Code** is a 24-character, alpha-numeric activation code which applies the appropriate license to your installation.
12. Click **Next** to continue.
13. If you are satisfied with the selections you have made and are ready to begin copying the Double-Take Availability for VMware Infrastructure files, click **Install**.
14. After the files have completed copying, click **Finish** to exit the installation program.

Installing Double-Take Availability automatically

The Double-Take Availability installation program can accept command-line parameters which allow you to automate the installation or upgrade process by running an unattended, or silent, installation. The automatic process allows you to pass parameters through to the installation program instead of entering information manually during the installation or upgrade.

Since the automated process does not prompt for settings, the settings are manually defined in a configuration file called DTSetup.ini. By default, DTSetup.ini contains two sections. The second section can be duplicated as many times as necessary. The first section, [Config], applies to any server not defined in the second (or duplicate of second) sections. The second (or duplicate of second) section, [MachineName], allows you to specify unique settings for individual servers. You have to modify the heading name (case-sensitive) to identify the server.

Review the following table to understand the different parameters available in DTSetup.ini.

DTSetupType

- **DTNT**—Both the Double-Take Availability server and client components will be installed.
- **DTCO**—Only the Double-Take Availability client components will be installed.
- **DTSO**—Only the Double-Take Availability server components will be installed.

If you are installing on Windows Server Core or Windows Hyper-V Server (standalone), the setup type will be server components only regardless of your setting.

DTActivationCode

A 24 character, alpha-numeric activation code which applies the appropriate license to the server. Multiple activation codes can be separated by a semi-colon.

DoubleTakeFolder

Any valid path specifying the location of the Double-Take Availability files

QMemoryBufferMax

Any integer representing the amount of system memory, in MB, to use for memory-based queuing

DiskQueueFolder

Any valid path to the location of the disk-based queue.

DiskQueueMaxSize

Any integer representing the amount of disk space, in MB, to use for disk-based queuing or the keyword **UNLIMITED** which will allow the queue usage to automatically expand whenever the available disk space expands

DiskFreeSpaceMin

Any integer representing the amount of disk space, in MB, that must remain free at all times

DTServiceStartup

- **Y** or **1**—Start the Double-Take service automatically
- **N** or **0**—Do not start the Double-Take service automatically

This parameter is not applied if your **DTSetupType** is DTCO.

Port

Any integer between 1024 and 65535 that identifies the Windows Firewall port used for Double-Take Availability.

Set_FWPORT

- **Y** or **1**—Set the Windows Firewall port for Double-Take Availability
- **N** or **0**—Do not set the Windows Firewall port for Double-Take Availability

---

Note: You must have Microsoft .NET installed on server before starting the automatic installation.

If you are using Windows 2008, but you are not using the built-in administrator account, Windows 2008 User Access Control will prompt you to confirm you want to install Double-Take Availability. To work around this issue, use the built-in administrator account when you are installing to each server. You may also disable User Access Control if that is acceptable for your environment.

Installing or upgrading automatically to a local machine

1. Create a temporary directory on the server. For example, create c:\temp_install.
2. On the CD, locate the files in a subdirectory under \setup\dtsw that is appropriate for your architecture, either i386, x64, or IA64. Copy the files from that subdirectory to the temporary directory.
3. From a command prompt, remove the read-only attributes from the files in the temporary directory by using the command **attrib *.* -r**.
4. Make a backup copy of the default DTSetup.ini file in the temporary directory.
5. Edit DTSetup.ini as needed using the values described in the previous table.
6. Determine the exact file name of your setup file by using the command dir setup*.* from the temporary directory command prompt. Depending on how you received your software (CD or web), your setup file name will be named setup.exe or setup_xxxx.exe where xxxx is four numbers that specify the build number. For example, your setup file might be called setup.exe or setup_1352.exe.
7. Run one of the following case-sensitive commands from the temporary directory, depending on if you have setup.exe or setup_xxxx.exe where xxxx is a four digit build number.

   setup /s /v"DTSETUPINI=\"c:\dtinstall\DTSetup.ini\" /qn"

   setup_xxxx /s /v"DTSETUPINI=\"c:\dtinstall\DTSetup.ini\" /qn"

Note: The command must be run from the directory where the temporary files are located as well as specifying that directory for the .ini file.

Spacing is critical with this command. A space should precede /s, /v, and /qn but should not appear anywhere else for the command to work correctly.

Installing or upgrading automatically to a remote machine

1. Create a temporary directory on the primary site server. For example, create z:\temp_install.
2. Share the temporary folder.
3. On the CD, locate the files in a subdirectory under \setup\dtsw that is appropriate for your architecture, either i386, x64, or IA64. Copy the files from that subdirectory to the temporary directory.
4. From a command prompt, remove the read-only attributes from the files in the temporary directory by using the command **attrib *.* -r**.
5. Make a backup copy of the default DTSetup.ini file in the shared folder.
6. Edit DTSetup.ini as needed using the values described in the previous table.
7. From each server where you want to install Double-Take Availability, map a drive to the temporary directory that you created in step 1. For example, you might map your m: drive to the share.
8. Determine the exact file name of your setup file by using the command dir setup*.* from the mapped drive command prompt. Depending on how you received your software (CD or web), your setup file name will be named setup.exe or setup_xxxx.exe where xxxx is four numbers that specify the build number. For example, your setup file might be called setup.exe or setup_1352.exe.
9. Run one of the following case-sensitive commands from the mapped drive, depending on if you have setup.exe or setup_xxxx.exe where xxxx is a four digit build number.

   setup /s /v"DTSETUPINI=\"m:\DTSetup.ini\" /qn"

   setup_xxxx /s /v"DTSETUPINI=\"m:\DTSetup.ini\" /qn"

Note: The command must be run from the shared folder as well as specifying that directory for the .ini file.

Substitute your mapped drive for m:\.

Spacing is critical with this command. A space should precede /s, /v, and /qn but should not appear anywhere else for the command to work correctly.

```
C:\>net use m: \\server_name\share
The command completed successfully
C:\>M:
M:\>setup_1352 /s /v"DTSETUPINI=\"m:\DTSetup.ini\" /qn"
```

Configuring your cluster for GeoCluster installation

If you want to use a GeoCluster configuration, where data is stored on volumes local to each node and replicated to each node in the cluster, complete the cluster configuration appropriate for the operating system you are using.

- Configuring your Windows 2003 cluster
- Configuring your Windows 2008 cluster

Configuring your Windows 2003 cluster

In a typical Windows 2003 MSCS shared disk cluster configuration, the quorum resource, by default, is the Local Quorum and is located on the first shared disk in the cluster. Because in a GeoCluster configuration there is no shared physical disk, the Local Quorum will not work as the quorum resource. You will need to choose one of the other Windows quorums. The recommended quorum resource for GeoCluster is the Majority Node Set or Majority Node Set with File Share Witness.

Note: If you are upgrading from a previous GeoCluster version and were using GeoCluster as a quorum, you must select another quorum type. GeoCluster can no longer be used as a quorum resource.

- **Local Quorum**—This quorum is for single node clusters and shared disk clusters. It cannot be used in a GeoCluster configuration.
- **Majority Node Set**—This quorum is for clusters with three or more nodes.
- **Majority Node Set with File Share Witness**—This quorum is for clusters with only two nodes. If you are using Windows 2003 Service Pack 1 or earlier, see the Microsoft support article 921181 for an update for the File Share Witness. If you are using Service Pack 2 or later, the update is not needed.

Use the following instructions as a guideline for configuring your Windows 2003 cluster. See your Windows cluster documentation as a complete reference.

1. Login with an account that has administrative rights on the domain and the local machine.
2. Create the cluster on the first node, if it is not already created. See your Windows documentation for instructions on how to create a cluster.
3. Add your additional nodes to the cluster. See your Windows documentation for instructions on how to add nodes to the cluster.
4. Install GeoCluster on each node of the cluster.
5. Configure your quorum. See your Windows documentation for instructions on configuring the quorum appropriate for your environment.
6. If desired, you can install GeoCluster on non-clustered client machines if you want to use Cluster Administrator to control the GeoCluster resources. Install GeoCluster, selecting the **Client Components Only** installation option.

Configuring your Windows 2008 cluster

The default quorum resource in a Windows 2008 environment will vary depending on your configuration (number of nodes, shared disks, and so on). The recommended quorum resource to use for GeoCluster is the Node and File Share Majority. There are other quorum types available. Review the following list to determine which quorum is appropriate for your environment.

- **Node Majority**—This quorum is recommended for clusters with an odd number of nodes. The cluster can handle failures of half of the nodes (rounding up) minus one and still stay online.
- **Node and Disk Majority**—This quorum is recommended for clusters with an even number of nodes. The cluster can handle failures of half of the nodes (rounding up), as long as the witness disk remains online, and still stay online. If the witness disk fails, the cluster can handle failures of only half of the nodes (rounding up) minus one and still stay online.
- **Node and File Share Majority**—This quorum is recommended for clusters with special configurations, such as GeoCluster. The cluster can handle failures of half of the nodes (rounding up), as long as the witness share remains online, and still stay online. If the witness share fails, the cluster can handle failures of only half of the nodes (rounding up) minus one and still stay online.
- **No Majority: Disk Only**—This quorum is not usually recommended. The cluster can handle failures of all nodes except one and still stay online.

Use the following instructions as a guideline for configuring your Windows 2008 cluster. See your Windows cluster documentation as a complete reference.

1. Login with an account that has administrative rights on the domain and the local machine.
2. Create the cluster, if it is not already created. See your Windows documentation for instructions on how to create a cluster.
3. Configure a Node and File Share Majority quorum. See your Windows documentation for instructions on how to configure the quorum.
4. If you are going to be using Hyper-V, install the Hyper-V server role on all nodes in the cluster. Make sure that you have the required Microsoft hotfixes applied, including KB958065 which is a failover clustering hotfix and KB950050.
5. Install GeoCluster on each node of the cluster.
6. If desired, you can install GeoCluster on non-clustered client machines if you want to use Cluster Administrator to control the GeoCluster resources. Install GeoCluster, selecting the **Client Components Only** installation option.
7. If you are going to be using Hyper-V, create your virtual machine from within Hyper-V. Be sure to leave the virtual machine off.

8. From Failover Cluster Management, create your application or service group.

Note: If you are creating a file server using clustered file shares, the path for the file share in the Failover Cluster Management wizard is case-sensitive. If the drive letter is uppercase, the path in the clustered file share wizard must also be uppercase. If the case does not match, the wizard will fail stating the path does not exist.

---

9. If you are using Hyper-V, add your virtual machine resource to the group. Any warnings about storage may be disregarded because the GeoCluster Replicate Disk will alleviate storage requirements.

Double-Take Console

The Double-Take Console is used to protect and monitor your servers and connections. Each time you open the Double-Take Console, you start at the **Home** page. This page provides a high-level overview of the status of your connections.

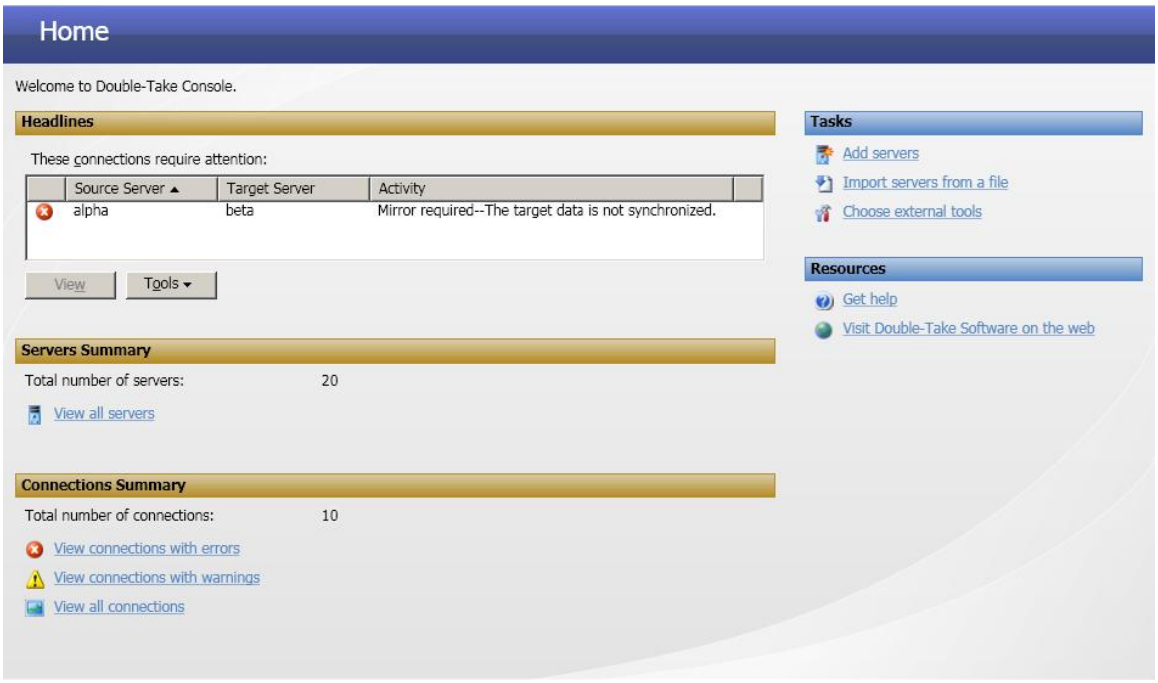

The appearance of the **Home** page is the same for all users. However, other console pages may have variances in the appearance or you may not see some pages at all. The pages and views depend on the Double-Take Software products that you have installed.

- Starting the console
- Getting started
- Managing servers
- Console options
- Other consoles

Starting the console

After you have installed the console, you can launch it by selecting **Start**, **Programs**, **Double-Take**, **Double-Take Console**.

Getting started

The first time you start the console, it is empty. In order to protect and monitor your servers, you must insert your servers in the console. You can insert servers manually, through Active Directory discovery, or from a console server configuration file.

Inserting servers manually

1. Select **Get Started** from the toolbar.
2. Select **Add servers** and click **Next**.
3. On the **Manual Entry** tab, specify the server information.
    - **Server**—This is the name of the server to be added to the console.
    - **User name**—Specify a user that is a member of the **Double-Take Admin** or **Double-Take Monitors** security group on the server.
    - **Password**—Specify the password associated with the **User name** you entered.
4. If necessary, specify the domain or protocol under **More Options**.
    - **Domain**—If you are working in a domain environment, specify the **Domain**.
    - **Protocol**—Specify the protocol type that the console will use to communicate with the server.
    - **Automatically detect protocol**—Double-Take Availability will check the server to determine the protocol type to use.
    - **XML web services protocol**—Select this option to use XML web services as your protocol. Select this option if your server is running Double-Take Availability version 5.2 or later.
    - **Legacy protocol**—Select this option to use the legacy proprietary Double-Take protocol. Select this option if your server is running Double-Take version 5.1 or earlier.
5. After you have specified the server information, click **Add**.
6. Repeat steps 3 through 5 for any other servers you want to add.
7. If you need to remove servers from the list of **Servers to be added**, highlight a server and click **Remove**. You can also remove all of them with the **Remove All** button.
8. When your list of **Servers to be added** is complete, click **OK**.

You will automatically be taken to the Manage Servers page.

Inserting servers through Active Directory discovery

1. Select **Get Started** from the toolbar.
2. Select **Add servers** and click **Next**.
3. Select the **Automatic Discovery** tab.
4. Click **Discover** to search Active Directory for servers running Double-Take Availability.
5. If you need to remove servers from the list of **Servers to be added**, highlight a server and click **Remove**. You can also remove all of them with the **Remove All** button.
6. When your list of **Servers to be added** is complete, click **OK**.

You will automatically be taken to the Manage Servers page.

Inserting servers from a server configuration file

You can share the console server configuration between machines that have the Double-Take Console installed. The console server configuration includes the server name, server communications ports, user name, encrypted password, and other internal processing information.

1. To export a server configuration file, select **File**, **Export**.
2. Specify a file name and click **Save**.

After the configuration file is exported, you can import it to another console.

When you are importing a console server configuration file from another console, you will not lose or overwrite any servers that already exist in the console. For example, if you have server alpha in your console and you insert a server configuration file that contains servers alpha and beta, only the server beta will be inserted.

1. To import a server configuration file, select **File**, **Import**.
2. Locate the console configuration file saved from the other machine and **Open**.

After the configuration file is imported, you will automatically be taken to the Manage Servers page.

Managing servers

To manage the servers in your console, select **Manage Servers** from the toolbar. The **Manage Servers** page allows you to view, add, edit, or remove servers from your console.

Column 1 (Blank)

The first blank column indicates the status of communications between the console and the server.

The console is attempting to communicate with the server.

The server is online and the console is communicating with it.

The server is offline, so the console cannot communicate with it.

The console has lost communication with the server.

Column 2 (Blank)

The second blank column indicates the security level

Administrator access

Monitor only access

No security access

Server

The name of the server

Activity

There are many different **Activity** messages that keep you informed of the server activity. Most of the activity messages are informational and do not require any administrator interaction. If you see error messages, check the server details.

Version

The product version information

Product

The products licensed for the server

Activation Code

The activation codes associated with the products licensed for the server

---

Add Servers 

Add a new server. This button leaves the **Manage Servers** page and opens the Add Servers page

View Server Details 

View detailed information about a server. This button leaves the **Manage Servers** page and opens the View Server Details page

Remove Server 

Remove the server from the console

Provide Credentials 

Change the login credentials for a server

View Server Events 

View event messages for a server. This button leaves the **Manage Servers** page and opens the View Server Events page.

Viewing server details

To view details about a specific server, select **View Server Details** from the toolbar on the **Manage Servers** page. The **View Server Details** page allows you to view details about a particular server.

Server name

The server name or IP address as added to the console

Status

There are many different **Status** messages that keep you informed of the server activity. Most of the status messages are informational and do not require any administrator interaction. If you see error messages, check the rest of the server details.

Activity

There are many different **Activity** messages that keep you informed of the server activity. Most of the activity messages are informational and do not require any administrator interaction. If you see error messages, check the rest of the server details.

Protocol

- **Automatically detect protocol**—The console automatically determined the protocol type to use to communicate with the server.
- **XML web services protocol**—XML web services is the protocol being used to communicate with the server. This server should be running Double-Take version 5.2 or later.
- **Legacy protocol**—Double-Take legacy, proprietary protocol is being used to communicate with the server. This server should be running Double-Take version 5.1 or earlier.

Port

The port used for communication with the server

Version

The product version information

Access

The security level granted to the specified user

User name

The user account used to access the server

Licensing

Licensing information for the server

- **Product**—The product associated with the license
- **Serial number**—The serial number associated with the license
- **Expiration date**—The date the license expires, if there is one
- **Code**—The activation code

Source connections

A list of any connections from this server. Double-clicking on a connection in this list will automatically open the View Connection Details page.

Target connections

A list of any connections to this server. Double-clicking on a connection in this list will automatically open the View Connection Details page.

Viewing server events

To view events associated with a specific server, select **View Server Events** from the toolbar on the **Manage Servers** page.The **View Server Events** page displays the same messages that are logged to the Windows Event Viewer. The list of events are displayed in the top pane of the page, although the Description is limited. When you highlight an event, the event details, including the full Description, are displayed in the bottom pane of the page.

- **Severity**—An icon or text that classifies the event, such as Error, Warning, Information, Success Audit, or Failure Audit
- **Time**—The date and time the event occurred
- **ID**—An identification number to help identify and track event messages.
- **Source**—The component that logged the event
- **Description**—The event details

You can filter the events displayed by using the **Filter** drop-down list or the **View warning events** and **View error events** toolbar buttons. To clear a filter, select **All events** in the **Filter** drop-down list. See Event messages for a complete list of the service and driver event messages.

Providing server credentials

To update the security credentials used for a specific server, select **Provide Credentials** from the toolbar on the **Manage Servers** page. When prompted, specify the **User name**, **Password**, and **Domain** of the account you want to use for this server. Click **OK** to save the changes.

Managing VirtualCenter servers

To manage your VirtualCenter servers, select **Go**, **Manage VirtualCenter Servers**. The **Manage VirtualCenter Server** page allows you to view, add, remove, or edit credentials for your VirtualCenter servers available in the console.

VirtualCenter Server

The name of the VirtualCenter server

Full Name

The full name of the VirtualCenter server

User Name

The user account being used to access the VirtualCenter server

---

Add VirtualCenter Servers 

Add a new VirtualCenter server. When prompted, specify the VirtualCenter server and a user account.

Remove Server 

Remove the VirtualCenter server from the console.

Provide Credentials 

Edit credentials for the selected VirtualCenter server. When prompted, specify a user account to access the VirtualCenter server.

Console options

There are several options that you can set that are specific to the Double-Take Console. To access these console options, select **Options** from the toolbar.

- Setting the frequency of console refreshes
- Setting the console communications port
- Updating the console software automatically

Setting the frequency of console refreshes

To access the console refresh rate, select **Options** from the toolbar. On the **Options** page, **Monitoring interval** specifies how often, in seconds, the console refreshes the monitoring data. The servers will be polled at the specified interval for information to refresh the console.

Setting the console communications port

To access the console communications port, select **Options** from the toolbar. On the **Options** page, **Default port** specifies the port that the console will use when sending and receiving data to Double-Take servers. By default, the port is 6325. Changes to the console port will not take effect until the console is restarted.

Note: If you are using an older Double-Take version, you will need to use the legacy protocol port. This applies to Double-Take versions 5.1 or earlier.

Updating the console software

By default, each time the console is started, it will automatically check the Double-Take Software web site to see if there is updated console software available. If there is updated console software available, an **Automatic Updates** section will appear on the **Home** page. Click **Get the latest update** to download and install the updated console software.

If you want to disable the automatic check for updates, click **Change automatic updates** or select **Options** from the toolbar. On the **Options** page, deselect **Automatically check for updates** to disable the automatic check.

You can also manually check for updates by selecting **Help**, **Check for Updates**.

- **Update available**—If there is an update available, click **Get Update**. The dialog box will close and your web browser will open to the Double-Take Software web site where you can download and install the update.
- **No update available**—If you are using the most recent console software, that will be indicated. Click **Close**.
- **No connection available**—If the console cannot contact the update server of if there is an error, the console will report that information. The console log contains a more detailed explanation of the error. Click **Check using Browser** if you want to open your browser to check for console software updates. You will need to use your browser if your Internet access is through a proxy server.

Other consoles

There are various Double-Take Availability consoles, many of which are being phased out over time. To help consolidate the consoles and help you locate the necessary workflows to complete your work, use the console called Double-Take Console.

To access this console, select **Start**, **Programs**, **Double-Take**, **Double-Take Console**. Select **Get Started** from the toolbar and then select the type of workload protection you want to establish. The appropriate workflow or console will then open.

Use the following links to find more information on the other consoles available.

- Replication Console
- Failover Control Center
- Full-Server Failover Manager
- Application Manager
- Double-Take Availability for VMware Infrastructure console

Replication Console

To open the Replication Console, select **Start**, **Programs**, **Double-Take**, **Double-Take Replication Console**.

From the Replication Console, you can manage, monitor, and control your data workload connections. The Replication Console is a two pane view. The views in the panes change depending on what is highlighted. For example, when the root of the tree in the left pane is selected, all of the machines in your environment running Double-Take Availability are displayed in the right pane. If you expand the tree in the left pane and select a server, any connections for that server are displayed in the right pane.

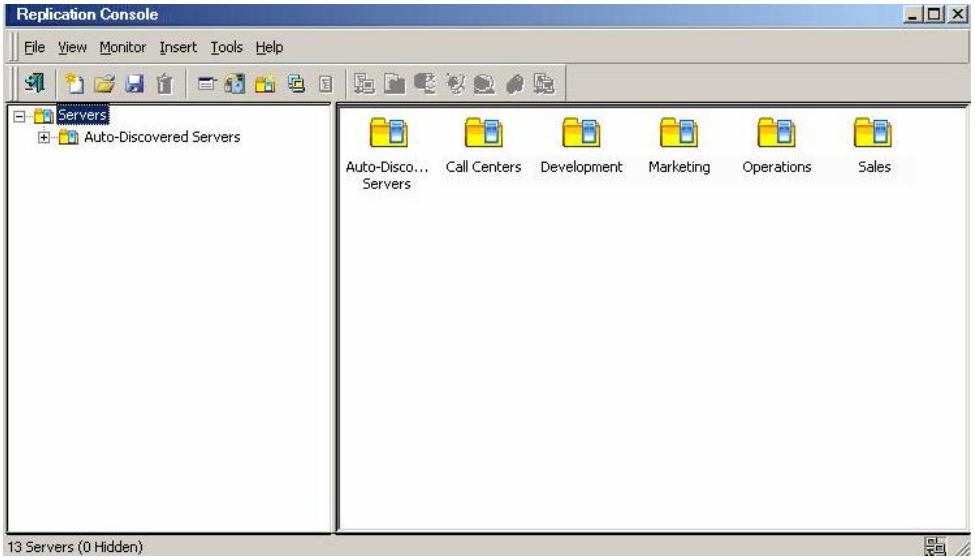

Note: You may not have access to some of the components or see certain display options if you are using a newer version of the Replication Console to control an older version of your source or target.

If you are logged in locally to the machine running the Replication Console, there will be no servers automatically populated in the Servers tree. You will have to manually insert each server.

- Logging on and off
- Managing the Replication Console tree
- Sharing group and server configuration
- Workspaces
- Clearing maintained security credentials

Logging on and off

To ensure protection of your data, Double-Take Availability offer multi-level security using native operating system security features. Privileges are granted through membership in user groups defined on each machine running Double-Take Availability. To gain access to a particular Double-Take Availability source or target, the user must provide a valid operating system user name and password and the specified user name must be a member of one of the Double-Take Availability security groups. Once a valid user name and password has been provided and the Double-Take Availability source or target has verified membership in one of the Double-Take Availability security groups, the user is granted appropriate access to the source or target and the corresponding features are enabled in the client. Access to Double-Take Availability is granted on one of the following three levels.

- **Administrator Access**—All Double-Take Availability features are available for that machine. For example, this access level includes creating replication sets and establishing Double-Take Availability connections.
- **Monitor Access**—Statistics can be viewed on that machine, but Double-Take Availability features are not available. For example, this access level does not allow the user to create or modify replication sets or create or modify Double-Take Availability connections, but does allow you to view the connection statistics for any established Double-Take Availability connections on that machine.
- **No Access**—The machine appears in the Double-Take Availability Replication Console and can be pinged using a scripting command, but no other access is available.

Use the following instructions when logging on and off of a server.

1. Highlight a machine on the left pane of the Replication Console. By double-clicking the machine name, Double-Take Availability automatically attempts to log you on to the selected machine using the ID that you are currently logged on with. Verify your access by the resulting icon.
2. If you have no access, the Logon dialog box will automatically appear. If you have monitor access or want to log on with a different username, right-click the machine name and select **Logon**.

3. Specify your **Username**, **Password**, **Domain**, and whether you want your password saved.
4. Click **OK** and verify your access by the resulting icon and log on again if necessary.

Note: When logging in, the user name, password, and domain are limited to 100 characters.

If your activation code is missing or invalid, you will be prompted to open the Server Properties **General** tab to add or correct the code. Select **Yes** to open the Server Properties dialog box or select **No** to continue without adding an activation code.

If the login does not complete within 30 seconds, it is automatically canceled. If this timeout is not long enough for your environment, you can increase it by adjusting the **Communication Timeout** on the **Configuration** tab of the Replication Console properties. Select **File**, **Options**, from the Replication Console to access this screen.

Double-Take Availability uses ICMP pings to verify server availability during the login process. If your Double-Take Availability server is across a router or firewall that has ICMP pings disabled, you will need to disable the Double-Take Availability ICMP ping verification. To do this, select **File**, **Options**, from the Replication Console and disable **Use ICMP to verify server availability**.

Double-Take Availability uses the current user's login credentials to attempt to log in to servers. This is called a unified login. If you want to disable unified logins, select **File**, **Options**, from the Replication Console and enable **Use Named Pipes for Login**.

Administrator rights 

This icon is a computer with a gear and it indicates the Double-Take Availability security is set to administrator access.

Monitor rights 

This icon is a computer with a magnifying glass and it indicates the Double-Take Availability security is set to monitor only access.

No rights 

This icon is a lock and it indicates the Double-Take Availability security is set to no access.

5. To log off of a Double-Take Availability machine, right-click the machine name on the left pane of the Replication Console and select **Logout**.

Managing the Replication Console tree

To better manage the servers that appear in the Replication Console, you can customize the server display to fit your needs. You can create groups and move servers to those groups to help you organize your environment. Within the groups, you can insert, remove, hide or unhide servers. Each of these functions is detailed in the following sections.

- **Groups**—The left pane of the Replication Console is a tree view of the Double-Take Availability servers. By default, the first group in the tree is the Discovered Servers group. All Double-Take Availability servers that are automatically discovered will be added to this group. Use server groups in a hierarchical structure to help you organize your environment.
  - Creating groups
  - Removing groups
- **Servers**—Within your server groups, you have the ability to further manage the servers that are displayed by using the following functions.
  - Moving servers
  - Inserting servers
  - Removing servers
  - Hiding servers
  - Unhiding servers

Creating groups

Use one of the following methods to create a new group:

- Right-click anywhere on the left pane of the Replication Console and select **New**, **Group**.
- Right-click on a group icon on the right pane of the Replication Console and select **New**, **Group**.
- Click **Add Group** on the toolbar.
- Use the menu bar and select **Insert**, **Group**.

The location of the new group that is created will depend on what was highlighted. If the root of the tree was highlighted, the new group will be created as a child of the root. If a group or server within a group was highlighted, the new group will be created as a child of that group. Name the newly inserted group with a unique name by typing over the default name and pressing Enter. This process is similar to naming a new folder in Windows Explorer. You can also rename an established group by double-clicking on the existing name. Type the new group name over the existing name and press Enter.

Removing groups

- Use one of the following methods to remove a group.
- Right-click on a group on the left of the Replication Console and select **Remove**.
- Right-click on a group on the right pane of the Replication Console and select **Remove**.
- Click **Remove Item** on the toolbar.
- Highlight a group and press the **Delete** key.

You will prompted to confirm the removal of the group and its subgroups and any servers (displayed or hidden) contained in them. Click **OK**.

If Active Directory discovery is enabled on the Replication Console, those servers that have Active Directory advertisement enabled will automatically be repopulated back in the default Discovered Servers group. If Active Directory discovery is disabled on the Replication Console or for individual servers, servers will need to be manually inserted into the Replication Console.

Note: The remove toolbar button also removes servers and replication sets, so make sure you have the correct item highlighted before clicking the toolbar button.

You cannot remove the default Discovered Servers group.

Moving Servers

Servers that are auto-populated can be moved to different groups within the Replication Console tree. You can move servers to groups by either of the following methods

- Drag and drop a server into the desired group. With this method you can move one server at a time in the left pane of the Replication Console. You can move more than one server at a time by using this method from the right pane of the Replication Console.
- Insert a server into the desired group. Inserting an existing server to the tree will move the first occurrence to that new location.

A Double-Take Availability server will only appear once within the entire Replication Console tree. Servers cannot be placed into multiple groups.

Inserting Servers

If a machine is not displayed on the Replication Console, it can be manually inserted. This feature is useful for machines that are across a router or on a different network segment.

Note: If a machine is manually inserted into the Replication Console, it will automatically be saved in your workspace and will appear the next time the Replication Console is started.

Use the following instructions to insert a server into the Replication Console.

1. A server that already exists in the tree will be moved to the currently selected group if you attempt to insert it again.
2. Right-click on a group and select **New**, **Server** or highlight a group and select **Insert**, **Server**.
3. Type the machine name or IP address and the port number, if it is different than the default.
4. Click **Test** to determine if the machine is running Double-Take Availability. At any time while Double-Take Availability is attempting to locate the machine, click **Stop** to cancel the test. If you do not manually test a machine before inserting it, Double-Take Availability will automatically test it for you.
5. Click **OK** to insert the server or Cancel if you do not want to insert that server.

Even if a machine is not running Double-Take Availability, you can still insert it in the Replication Console.

Removing Servers

If you do not want to see a server in the Replication Console, it can be permanently removed from the display. You might need to remove a server that was manually added, if that server is no longer needed. Or if there are servers within the network that another administrator is responsible for, you can remove them from your display.

If a server is listed in Active Directory and Active Directory discovery is enabled, a removed server will automatically be added back to the server list .

To remove a server, right-click on the server in the left or right pane of the Replication Console and select **Remove**. You can also select **Remove Item** from the toolbar.

If Active Directory discovery is enabled on the Replication Console, those servers that have Active Directory advertisement enabled will automatically be repopulated back in the default Discovered Servers group. If Active Directory discovery is disabled on the Replication Console or for individual servers, servers will need to be manually inserted into the Replication Console.

Note: The remove toolbar button also removes servers and replication sets, so make sure you have the correct item highlighted before clicking the toolbar button.

You cannot remove the default Discovered Servers group.

Hiding Servers

If you do not want to see a server in the Replication Console and do not want to disable Active Directory discovery, you can hide the server from view. This keeps the server in the Replication Console's internal list of servers, but does not display it in the server tree, any dialog boxes, or any field/menu selections.

To hide a server, right-click on a server in the left or right pane of the Replication Console and select **Hide**.

Note: If you attempt to insert a server that is already in the tree but hidden, you will be prompted to unhide the server and insert it into the selected group.

Be careful if you hide a server with an established Double-Take Availability connection. If that connection goes into an error state, you will not be able to see the connection in the Replication Console. The Double-Take Availability log, Event Viewer, and other monitoring methods will still be functioning to alert you to the error. Hiding the server only removes it from the Replication Console display.

If a target server with an established Double-Take Availability connection is hidden and you open the Connection Manager for that connection via the source, you will see the target and IP address displayed in the **Target** and **Route** fields, respectively. This is the only time you will see a hidden server.

Unhiding Servers

You can unhide, or display, a hidden server at any time you want to access that server. The server will be displayed in the server tree, all dialog boxes, and field/menu selections.

To unhide, or display, a hidden server, you can insert the server or use the following instructions.

1. Select **View**, **Unhide Servers**.
2. Select one or more servers by using Ctrl-click or Shift-click. You can also click **Select All** to select all of the servers in the list.
3. Click **Unhide**.

Before moving a group that contains at least one subgroup with at least two hidden servers, you must unhide all of the servers. After the servers have been unhidden, move the group and then hide the servers again. Any attempt to move a group containing subgroups with hidden servers will result in an error when saving the workgroup or exiting the Replication Console.

Sharing group and server configuration

All of the group and server information is stored on the local machine for each user. When you close the Replication Console, the group information is saved and will be persisted the next time you open the Replication Console on this machine. If you want to share the group configuration with another user or machine, you can export the group configuration information (**File**, **Export server group configuration**) to an .xml file. That file can then be copied and imported (**File**, **Import server group configuration**) by another user on this machine or to another machine.

Workspaces

The Replication Console workspace contains the display of the panes of the Replication Console and any servers that may have been inserted. Multiple workspaces can be used to help organize your environment or to view settings from another machine.

- Saving a workspace
- Opening a workspace

Saving a workspace

As you size, add, or remove windows in the Replication Console, you can save the workspace to use later or use on another Double-Take Availability client machine. Select **File** and one of the following options.

- **Save Workspace**—Save the current workspace. If you have not previously saved this workspace, you must specify a name for this workspace.
- **Save Workspace As**—Prompt for a new name when saving the current workspace.

Opening a workspace

From the Replication Console, you can open a new workspace or open a previously saved workspace. Select **File** and one of the following options.

- **New Workspace**—Open an untitled workspace with the default Double-Take Availability window settings.
- **Open Workspace**—Open a previously saved workspace.

Clearing maintained security credentials

To remove cached credentials, select **File**, **Options** and select the **Security** tab. To remove the security credentials, enable **Clear Cached Security Credentials** and then click **OK**.

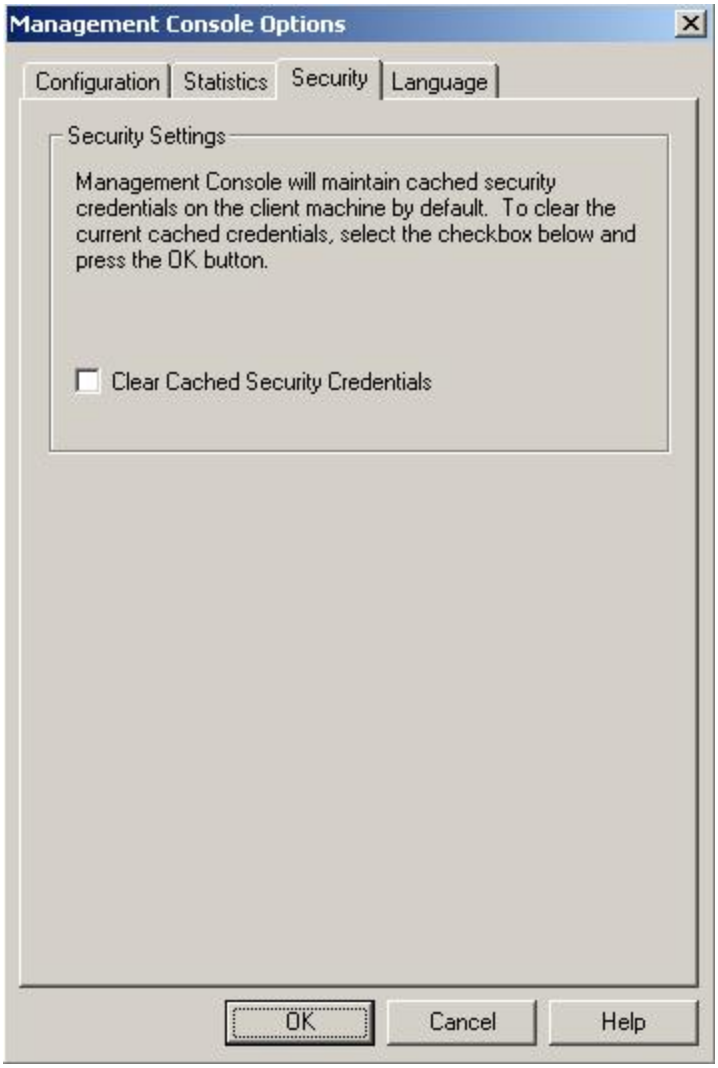

Failover Control Center

To open the Failover Control Center, select **Start**, **Programs**, **Double-Take**, **Availability**, **Double-Take Availability Failover Control Center**.

The Failover Control Center should be run from your target server or a client machine. Do not run the Failover Control Center from your source server.

From the Failover Control Center, you can manage, monitor, and control failover for your Double-Take Availability servers. The Failover Control Center displays a main window for monitoring failover activity. Control buttons to the right allow you to configure and manage your servers.

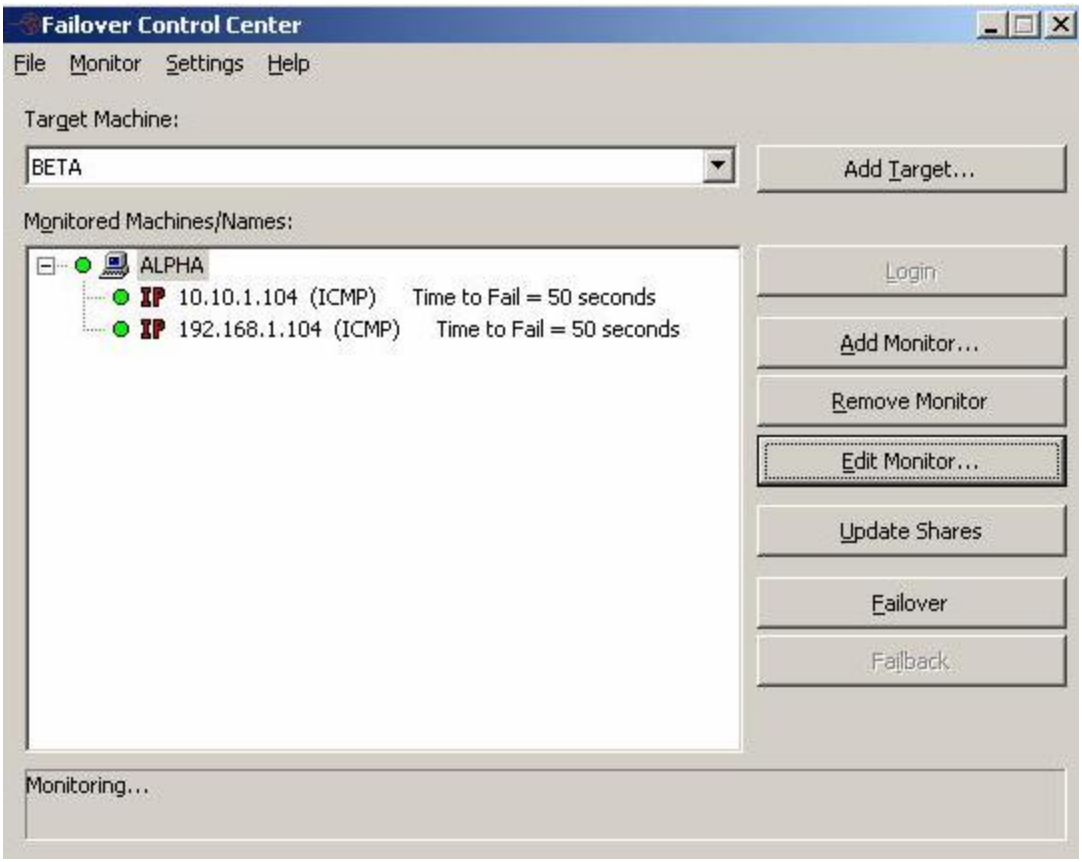

- Configuring communication ports
- Configuring the console refresh rate
- Clearing maintained security credentials

Configuring communication ports

The Failover Control Center uses port 6320, by default for Double-Take Availability communications. To view or modify the port settings in the Failover Control Center, select **Settings**, **Communications**.

Configuring the console refresh rate

The failover client periodically requests information from the source and target. Depending on the type of information, the request may be a machine-specific request, like obtaining the **Time to Fail** status from a target, or may be a general request, like determining which machines are running Double-Take Availability.

The rate at which these requests are made can be modified through the Failover Control Center refresh rate dialog box. Select **Settings**, **Refresh Rate**. The default update interval is one second. A lower refresh rate value updates the information in the Failover Control Center window's **Monitored Machines** tree more often, but also generates more network traffic and higher utilization on the client and target machines. A higher refresh rate value updates the information less frequently, but minimizes the network traffic.

Clearing maintained security credentials

To remove cached credentials, access the credentials security option, by selecting **Settings**, **Security**. To remove the security credentials, enable **Clear Cached Security Credentials** and then click **OK**.

Full-Server Failover Manager

To open the Full-Server Failover Manager, select **Start**, **Programs**, **Double-Take**, **Availability**, **Double-Take Availability Full-Server Failover Manager**.

The Full-Server Failover Manager allows you to create your source and target connection, monitor your full-server workload protection, manage your full-server snapshots, and initiate full-server failover.

The client is a simple interface with four numbered steps. Steps 1 and 2 for the source and target have to be completed before the other steps are available or the **Protection Status** is displayed.

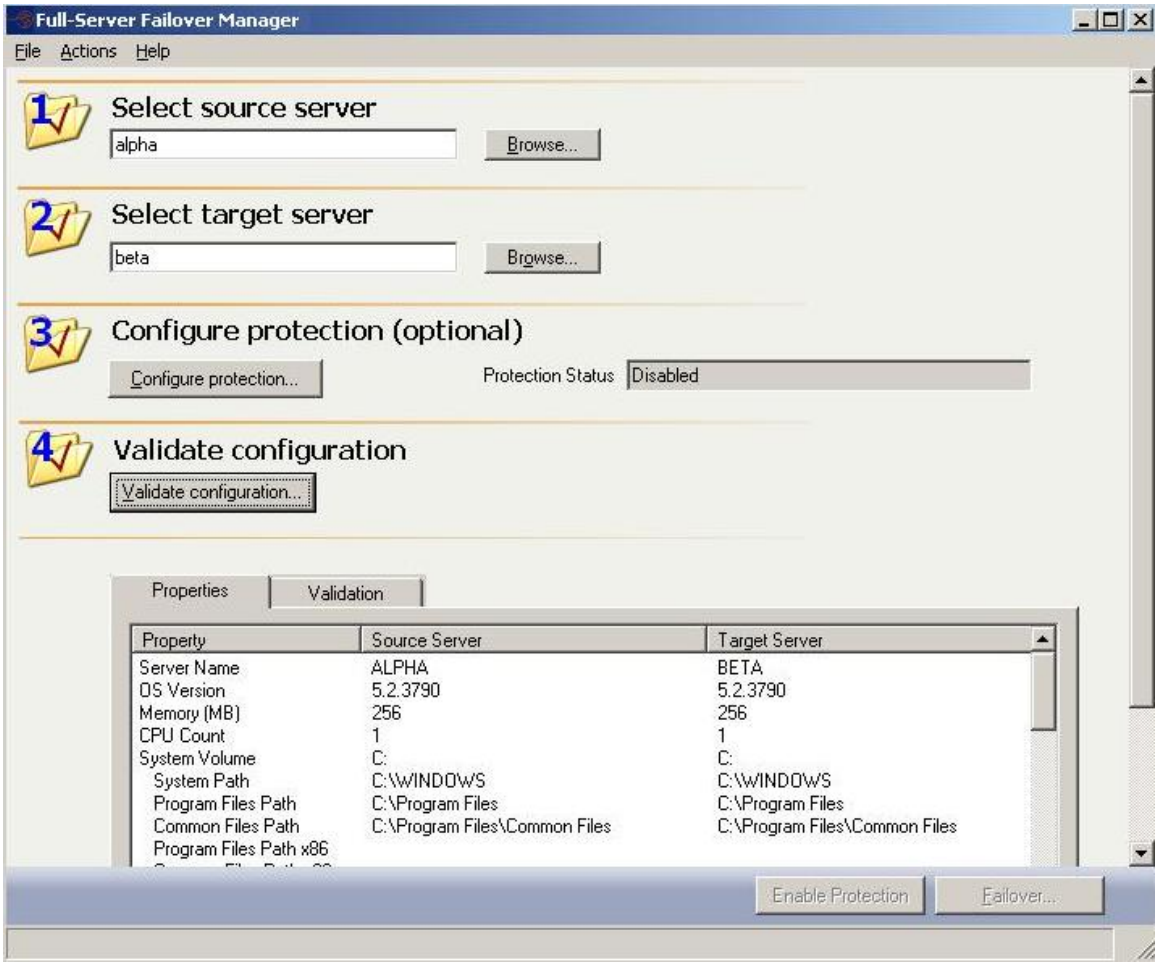

- Configuring the console refresh rate
- Configuring the level of detail to log
- Clearing maintained security credentials
- Configuring the monitoring method for server availability
- Saving and reusing configuration options

Configuring the console refresh rate

Select **File**, **Options**. By default, the main window of Full-Server Failover Manager will automatically update every five (5) seconds. If desired, you can modify the **Refresh Interval**. You can also refresh the main window manually by selecting **File**, **Refresh**.

Configuring the level of detail to log

Select **File**, **Options**. By default, Full-Server Failover Manager creates a log that maintains all processing information. If desired, you can disable the option **Enable Verbose Logging** so that only basic processing information is logged.

Clearing maintained security credentials

Select **File**, **Options**. By default, Full-Server Failover Manager will save the user credentials supplied for each servers. If desired, you can enable **Clear cached credentials on exit** so they are not saved. You can also clear the user credentials manually by selecting **File**, **Clear Cached Credentials**.

Configuring the monitoring method for server availability

Select **File**, **Options**. By default, Full-Server Failover Manager will use ICMP pings to check for server availability. If you do not want to use ICMP pings, select **Do not use ICMP Ping**. When this option is selected, the Full-Server Manager will use the Double-Take service to check for server availability.

Saving and reusing configuration options

After you have created a protection pair and configured any of the optional settings you can save those settings so that you can reuse them for future pairs of servers. Once you have the settings the way you want them, save them by selecting **File**, **Defaults**, **Save Current Settings as Defaults**. This creates a file called FFMDefaults.xml, which will automatically be the default settings the next time you use the Full-Server Failover Manager. If desired, you can rename the FFMDefaults.xml file and then save a new set of defaults to FFMDefaults.xml to be used by the Full-Server Failover Manager. This would allow you to have multiple failover configurations, which can be more easily interchanged. You can also use these different files to initiate failovers without using the Full-Server Failover Manager GUI.

Note: Because network adapters are uniquely identified on each server, the **Network Mapping**is not stored in the default settings.

If you want to reset the configuration settings back to the default settings, select **File**, **Defaults**, **Reset Defaults**.

Application Manager

To open the Application Manager, select **Start**, **Programs**, **Double-Take**, **Availability**, **Double-Take Availability Application Manager**.

Note: You can also start the Application Manager from the command line.

- To start it in standard mode, run the command dtam /application, where /application is /exchange, /sql, /fileprint, /blackberry, or /sharepoint.
- To start it in advanced mode, run the command dtam /application /advanced, where /application is /exchange, /sql, /fileprint, /blackberry, or /sharepoint.

The Application Manager allows you to establish application protection, monitor that protection, and initiate application failover and failback.

When you select an application to protect in the **Tasks** list on the left pane, the **Setup** tab on the right pane is a simple interface with four numbered steps. Steps 1 and 2 are for the domain and servers. Step 3 is optional configuration and step 4 validates the servers.

After protection has been established, use the **Monitor** tab to check on the status of your protection.

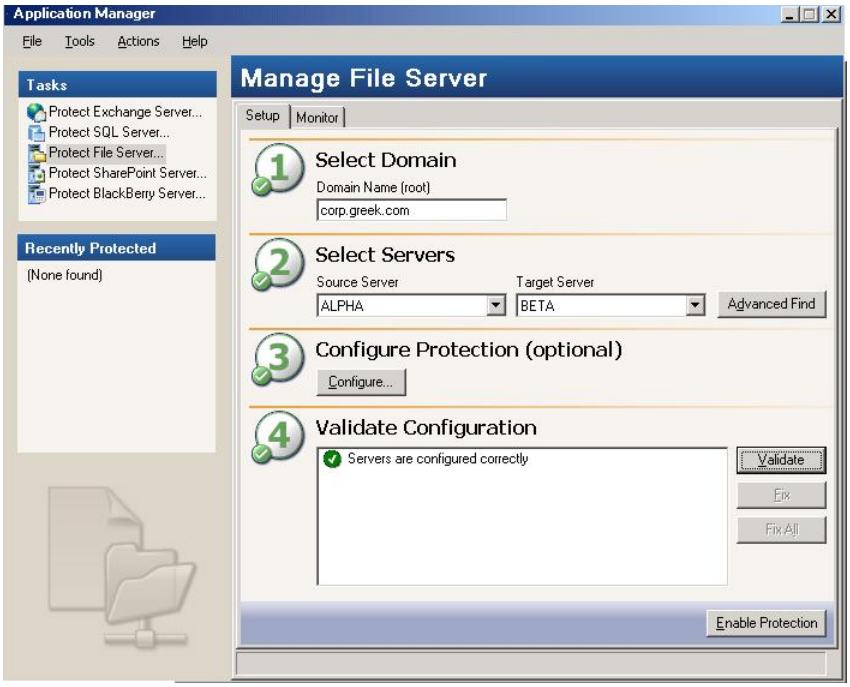

- Adding or managing servers
- Changing Application Manager preferences

Adding or managing servers

Step 2 of the application protection workflow is to select your source and target servers. If no servers are populated in the lists (perhaps the server you need is in a child domain), click **Advanced Find** to add servers to the lists. **Advanced Find** is not available for all application protections.

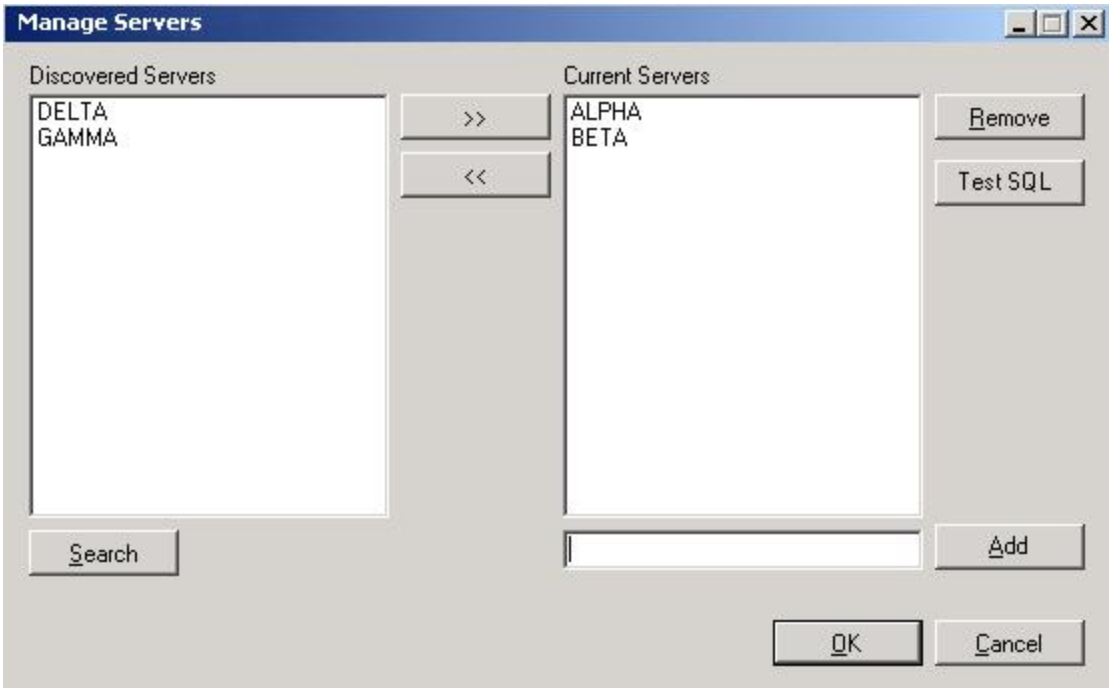

1. Click **Search** to locate all the application servers that Application Manager can discover in the domain. If you have a large number of servers in Active Directory, this search may take awhile.
2. Highlight servers in the **Discovered Servers** list and move them to the **Current Servers** list to add them to the Application Manager console.
3. To add a non-discovered server, type the server name below the **Current Servers** list and click **Add**.
4. To remove any servers from the Application Manager console, highlight the server name in the **Current Servers** list and click **Remove**.
5. For some applications, you can click the **Test SQL** button to have Application Manager check to see if the application is installed and accessible on the selected server.
6. When you have finished managing your servers, click **OK**.

Changing Application Manager options

To modify Application Manager options, select **Tools, Options** and modify any of the following options.

- **Service Listen Port**—This is the port used for Double-Take Availability communications. This port must be the same on the client machine and the source and target servers.
- **Enable automatic adjust of refresh interval**—When this option is enabled, the Application Manager will automatically adjust the rate at which protection status is refreshed.
- **Refresh Interval**—If you want to specify an exact interval for refreshing the Application Manager protection status, disable the automatic adjustment option and specify a length of time, in seconds, for to refresh the status.
- **Maximum log file size**—Specify the maximum size of the dtam.verbose.log file. When the maximum size is reached, the log file will be renamed to dtam.verbose.prev.log and a new log file will be used.
- **Enable verbose logging**—When enabled, this options logs all user interactions in the Application Manager to the dtam.verbose.log file.
- **Always show protection details**—When enabled, the **Protection Details** area on the **Monitor** tab will be expanded by default. When disabled, the area will be collapsed by default.
- **Display statistics values in bytes**—When enabled, values displayed in the **Protection Details** area on the **Monitor** tab will be shown in bytes. When disabled, the values will be shown in MB, GB, or TB.
- **Load last selected server upon startup**—This option automatically reconnects to the last protected source and target pair each time the Application Manager is started.
- **Enable Alternative DNS**—This option opens the Application Manager in \altdns mode on subsequent restarts so that Application Manager will not check for Microsoft DNS. See Non-Microsoft DNS for more information.
- **Display advanced options**—This option opens the Application Manager in \advanced mode on subsequent restarts.
- **Use Primary Dns Zone**—When enabled, the next time you start Application Manager, the server's primary DNS suffix will be used for the fully-qualified domain name if the server is listed in more than one DNS zone.
- **Clear Cached Credentials**—Click this button to clear all server credentials stored in Application Manager.

Double-Take Availability for VMware Infrastructure console

To open the Double-Take Availability for VMware Infrastructure console, select **Start**, **Programs**, **Double-Take**, **Availability**, **Double-Take Availability for VMware Infrastructure**.

The first time you use the console or if you have not saved your login information, you will be prompted to provide login information. Specify the **Server**, which is the machine running the Double-Take Availability for VMware Infrastructure service, and a **User name** and **Password**. If you do not want to provide login information each time you open the console, enable **Save DTAVI connection information**.

The Double-Take Availability for VMware Infrastructure console allows you to establish protection of host-level virtual disk files (the .vmdk files) from an ESX source to an ESX target. You can also initiate failover and failback.

The left pane is a tasks-style pane. When an item in the left pane is selected, the right pane of the console display updates to the corresponding workflow or page.

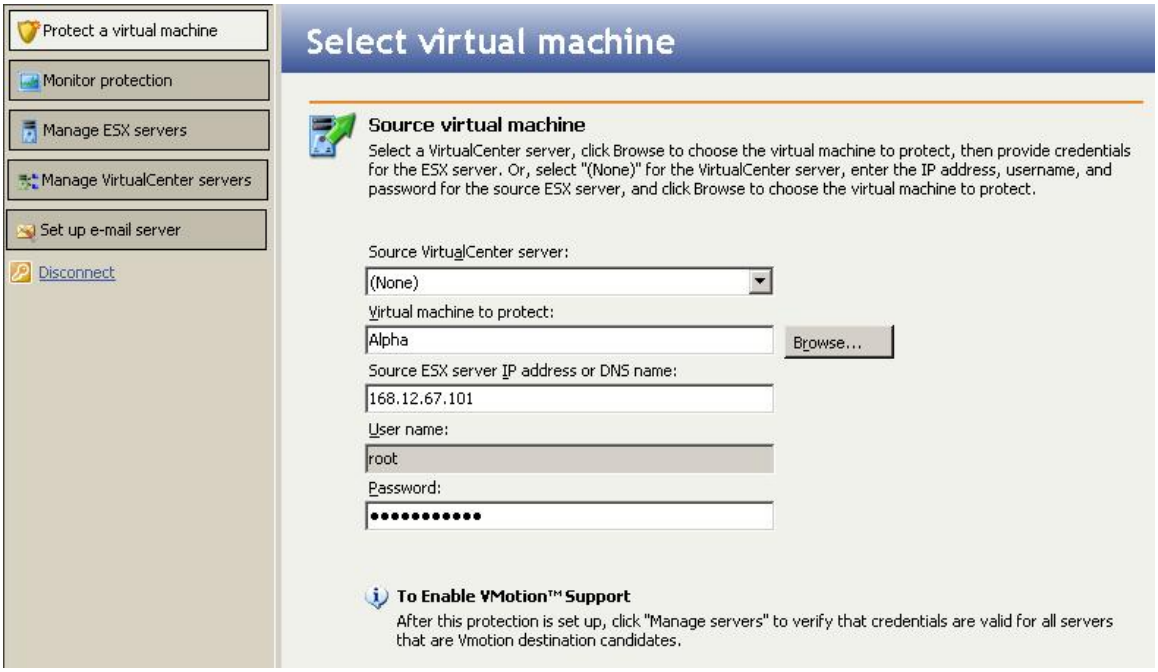

- Managing activation codes
- Managing VirtualCenter servers
- Managing ESX servers
- Setting up an e-mail server

Managing activation codes

You can manage your Double-Take Availability for VMware Infrastructure activation codes by selecting **Go**, **Manage activation codes**.

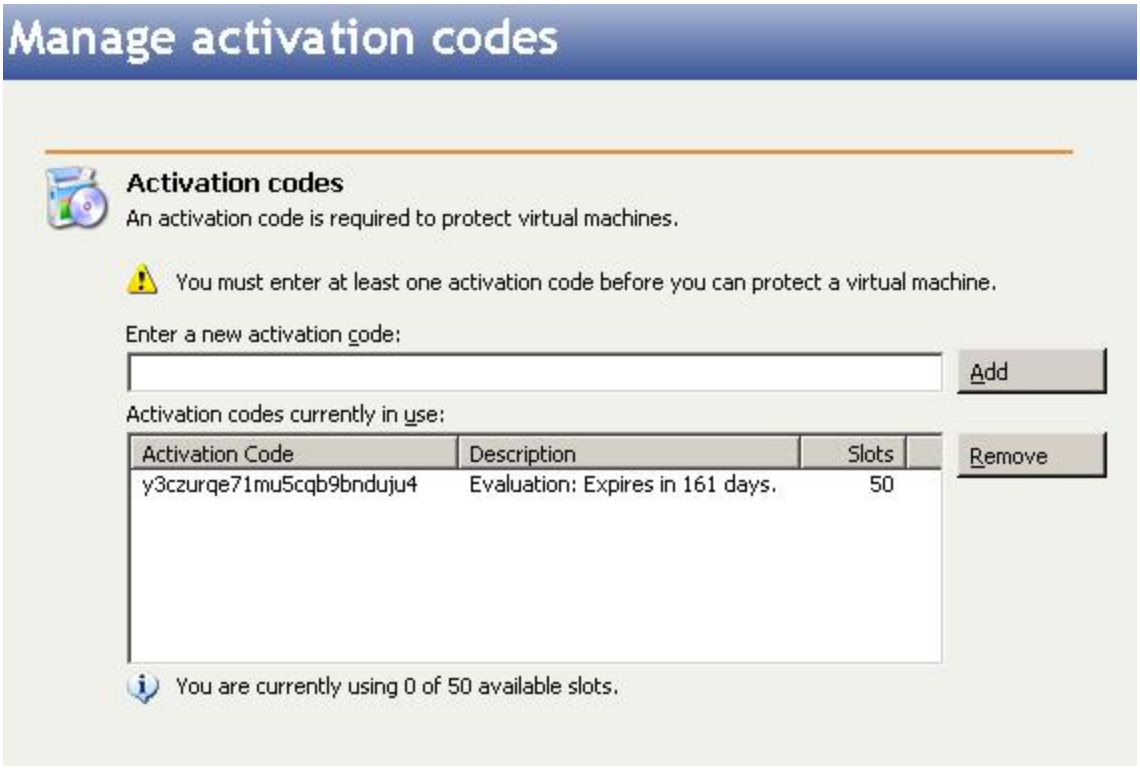

Enter a new activation code and click **Add**. To remove a code, highlight in the list and click **Remove**.

Each activation code corresponds to a number of slots, where each slot represents the capacity to protect a single virtual machine in your environment. Each time protection is established, Double-Take Availability for VMware Infrastructure will update the available number of slots for subsequent protections.

Note: If you are using VMware bundle licensing, each slot represents an ESX server, rather than a protection. Therefore, entering an ESX host by IP address and again by DNS name will cause a duplicate entry using an additional license slot. Therefore, when you add ESX servers, enter either the IP address or the DNS name but not both.

Managing VirtualCenter servers

To manage your VirtualCenter servers, select **Manage VirtualCenter servers** from the left pane of the console.

- **Adding a VirtualCenter server**—Click **Add VirtualCenter server** on the toolbar. On the **Add VirtualCenter server** page, specify the **IP address or DNS Name** of the VirtualCenter server and supply a **User name** and **Password**. Click **Save** to insert the VirtualCenter server.
- **Configuring credentials for a VirtualCenter server**—Highlight a VirtualCenter server in the list and click **Configure VirtualCenter server** on the toolbar. On the **Set VirtualCenter server credentials** page, specify the updated **User name** and **Password** and click **Save**.
- **Removing a VirtualCenter server**—Highlight a VirtualCenter server in the list and click **Remove VirtualCenter server** on the toolbar.

Managing ESX servers

To manage your ESX servers, select **Manage ESX servers** from the left pane of the console. Double-Take Availability for VMware Infrastructure scans to find ESX servers that are VMotion destination candidates, based upon SAN connectivity. The **Credentials Cached** column in the table identifies servers that need to have credentials added. To add the credentials, highlight a server in the list and click **Configure ESX server** on the toolbar. On the **Configure ESX server** page, add, edit, or remove a user. If prompted, specify the password associated with the user. Click **Done** to save the modficiations.

If you need to add an ESX server, click **Add ESX server** on the toolbar. On the **Add ESX server** page, specify the **VirtualCenter server**, the **IP address or DNS name** of the ESX server, a **User name**. and **Password**. Click **Save** to insert the ESX server.

If you need to remove an ESX server, highlight an ESX server in the list and click **Remove ESX server**.

Setting up an e-mail server

To set up an e-mail server, select **Set up e-mail server** from the left pane of the console. E-mail configuration applies to all protection jobs. Specify the e-mail server configuration.

- **From address**—Specify the e-mail address that you wan to appear in the From field of each message.
- **SMTP server**—Specify the SMTP server using the full Active Directory DNS name, the IP address, or the NetBIOS short name.
- **User name**—Specify a user account with privileges to send e-mail messages from your SMTP server.
- **Password**—Specify the password associated with the User name you entered.

Click **Save** to save your settings.

Workload protection

Double-Take Availability flexible configurations allow you to protect different workloads depending on the needs of your organization.

- Specific data—You can protect specific data (volumes, directories, files, and/or wildcards). In the event of a failure, the data you protected is available on the target.
- Entire server—You can protect an entire server, including the data and system state, which is the server's configured operating system and applications. In the event of failure, the target becomes the source.
- Applications—You can protect applications, including Exchange, SQL, SharePoint, BlackBerry, or a Windows file server. In the event of a failure, the data you protected is available on the target.
- Virtual servers—You can protect physical or virtual machine to a virtual machine on the VMware ESX or the Microsoft Hyper-V platform.
- Clusters—You can protect your cluster configuration, either a standard cluster or a GeoCluster.

Data protection

Protecting specific data consists of two main tasks - creating a replication set (to identify the data to protect) and connecting that replication set to a target.

You have the following data protection options.

- **Automated process**—If you would like to use an automated process that walks you through both the replication and connection tasks, you only need to complete the steps Establishing a connection using the automated Connection Wizard.
- **Manual process**—If you want to go through the tasks manually, begin by Creating a replication set and then continue with Establishing a connection manually using the Connection Manager.
- **NAT or firewall**—If your environment has a NAT or firewall configuration, you will need to begin with Creating a replication set and then follow the instructions for Establishing a connection across a NAT or firewall.
- **Simulating a connection**—If you want to simulate a connection for planning purposes, begin by Creating a replication set and then continue with Simulating a connection.

Once your connection is established, move on with Data workload failover to ensure high availability.

Establishing a connection using the automated Connection Wizard

The Connection Wizard guides you through the process of protecting your data. It helps you select a source, identify the data from your source that will be included in the replication set, and select a target.

1. From the Replication Console, start the Connection Wizard to establish your connection. If you have just opened the Replication Console, you can click **Make a connection** from the right pane of the Replication Console. If the quick launch screen is no longer visible, select **Tools**, **Connection Wizard**.

   Note: If the Servers root is highlighted in the left pane of the Replication Console, the **Connection Wizard** menu option will not be available. To access the menu, expand the server tree in the left pane, and highlight a server in the tree.

   ---

2. The Connection Wizard opens to the Welcome screen. Review this screen and click **Next** to continue.

   Note: At any time while using the Connection Wizard, click **Back** to return to previous screens and review your selections.

   ---

3. If you highlighted a source in the Replication Console, the source will already be selected. If it is not, select the Double-Take Availability source. This is the server where the files reside that you want to protect.

   Note: Double-Take Availability will automatically attempt to log on to the selected source using the identification of the user logged on to the local machine. If the logon is not successful, the Logon dialog box will appear prompting for your security identification. When logging in, the user name, password, and domain are limited to 100 characters.

   ---

4. Click **Next** to continue.
5. If you highlighted a target in the Replication Console, the target will already be

selected. If it is not, select the Double-Take Availability target. This is your backup server that will protect the source.

Note: Double-Take Availability will automatically attempt to log on to the selected target using the identification of the user logged on to the local machine. If the logon is not successful, the Logon dialog box will appear prompting for your security identification. When logging in, the user name, password, and domain are limited to 100 characters.

---

6. Click **Next** to continue.
7. Choose to create a new replication set or use a replication set that already exists.
   - **Create a new replication set with this name**—If you choose to create a new replication, specify a replication set name.
   - **Use this replication set**—If you choose to use an existing replication set, specify the name of that replication set by selecting it from the pull-down menu.
8. Click **Next** to continue.
9. If you are creating a new replication set, a tree display appears identifying the volumes and directories available on your selected source server. Mark the check box of the volumes and/or directories you want to protect and click **Next** to continue.
10. Select the location on the target where the data will be stored.

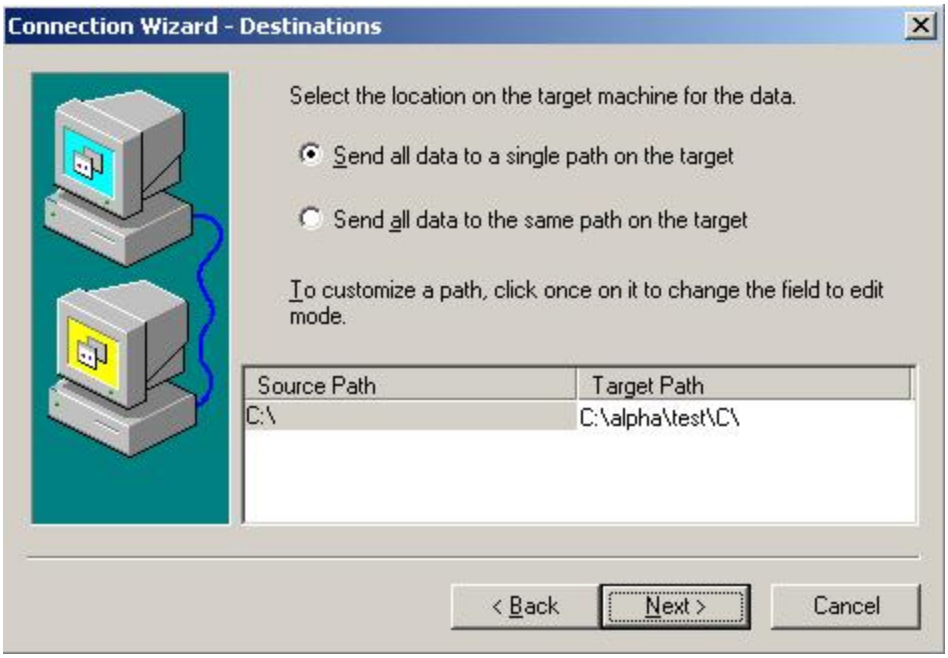

- **Send all data to a single path on the target**—This option sends all selected volumes and directories to the same location on the target. The default location is \source_name\replication_set_name\volume_letter.
- **Send all data to the same path on the target**—This option sends all selected volumes and directories to the same directories on the target. For example, c:\data and d:\files on the source will go to c:\data and d:\files on the target.
- **Custom**—To select a custom path, click once in the **Target Path** field and modify the drive and directory to the desired location.

11. Click **Next** to continue.
12. Review your selections on the summary screen. If your Connection Wizard settings are correct, establish your connection by completing one of the following two options.
    - If you do not want to set advanced options, click **Finish**. The Connection Wizard will close, the connection will be established, and mirroring and replication will begin.
    - If you want to set advanced options, click **Advanced Options**. The Connection Wizard will close and the Double-Take Availability Connection Manager will open. The **Servers** tab will be completed.

Creating a replication set

Before you can establish a connection, you must create a replication set.

1. From the Replication Console, highlight a source in the left pane of the Replication Console and select **Insert**, **Replication Set** from the menu bar. You can also right-click on the source name and select **New**, **Replication Set**.
2. A replication set icon appears in the left pane under the source. By default, it is named New Replication Set. Rename the newly inserted replication set with a unique name by typing over the default name and pressing **Enter**. This process is similar to naming a new folder in Windows Explorer.
3. Expand the tree under the replication set name to view the volume and directory tree for the source.

Note: The default number of files that are listed in the right pane of the Replication Console is 2500, but this is user configurable. A larger number of file listings allows you to see more files in the Replication Console, but results in a slower display rate. A smaller number of file listings displays faster, but may not show all files contained in the directory. To change the number of files displayed, select **File**, **Options** and adjust the **File Listings** slider bar to the desired number.

To hide offline files, such as those generated by snapshot applications, select **File**, **Options** and disable **Display Offline Files**. Offline files and folders are denoted by the arrow over the lower left corner of the folder or file icon.

---

4. Identify the data on the source that you want to protect by selecting volumes, drives, directories, and/or specific files.

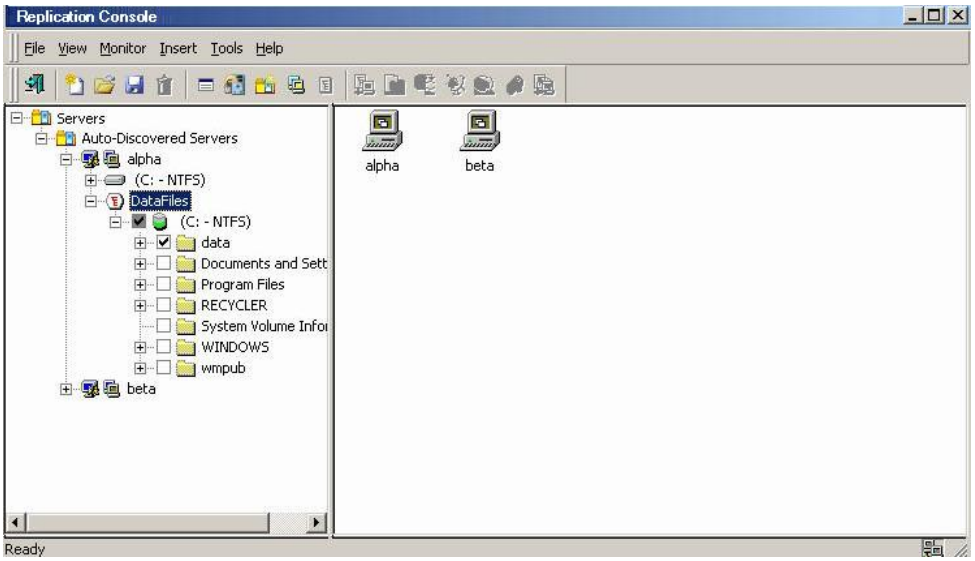

Note: Be sure and verify what files can be included by reviewing Replication capabilities.

5. After selecting the data for this replication set, right-click the new replication set icon and select **Save**. A saved replication set icon will change from red to black.

Establishing a connection manually using the Connection Manager

After you have created a replication set, you can establish a connection through the Connection Manager by connecting the replication set to a target.

1. From the Replication Console, open the Connection Manager to establish the connection.
   - Highlight the replication set and select **Tools**, **Connection Manager**.
   - Right-click on the replication set and select **Connection Manager**.
   - Drag and drop the replication set onto a target. The target icon could be in the left or right pane of the Replication Console.
2. The Connection Manager opens to the **Servers** tab. Depending on how you opened the Connection Manager, some entries on the **Servers** tab will be completed already. For example, if you accessed the Connection Manager by right-clicking on a replication set, the name of the replication set will be displayed in the Connection Manager. Verify or complete the fields on the **Servers** tab.

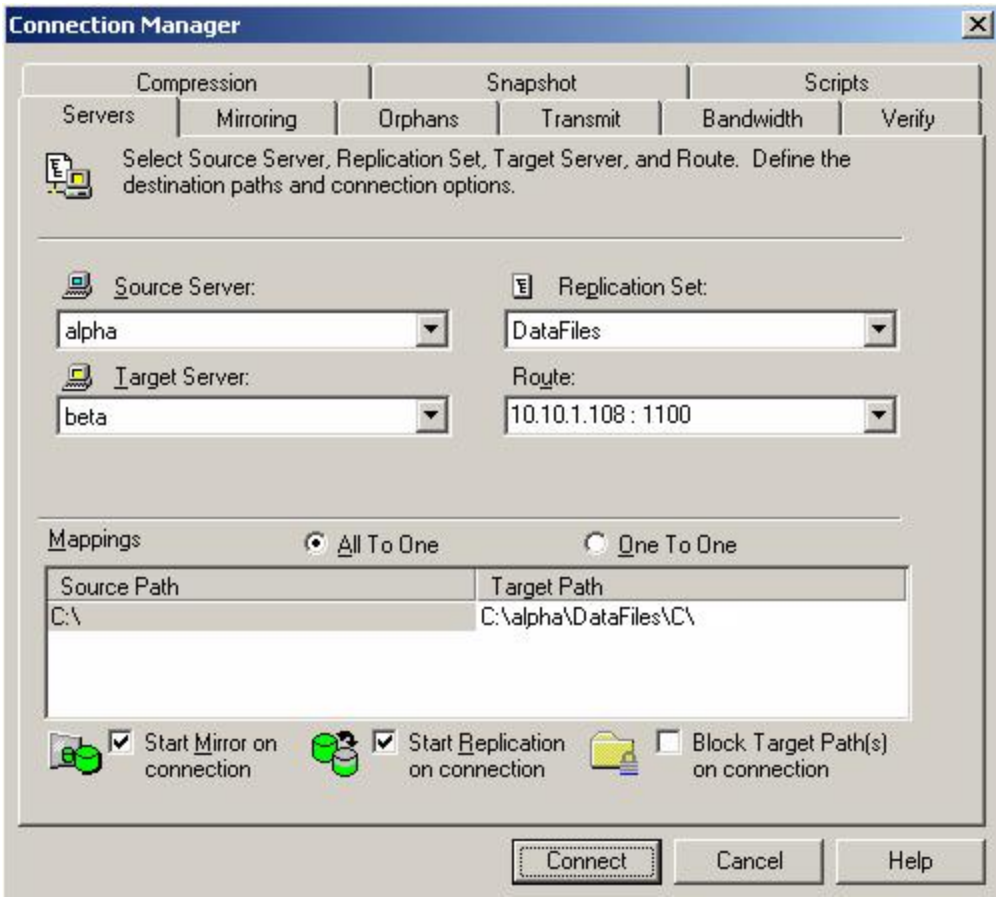

- **Source Server**—Specify the source server that contains the replication set that is going to be transmitted to the Double-Take Availability target.
- **Replication Set**—At least one replication set must exist on the source before establishing a connection. Specify the replication set that will be connected to the target.
- **Target Server**—Specify which Double-Take Availability target will maintain the copy of the source's replication set data. You can specify a machine name, IP address, or virtual IP address.
- **Route**—This is an optional setting allowing you to specify the IP address and port on the target that the data will be transmitted through. This allows you to select a different route for Double-Take Availability traffic. For example, you can separate regular network traffic and Double-Take Availability traffic on a machine with multiple IP addresses.
- **Mappings**—You must specify the location on the target where the source's replication set data will reside. Double-Take Availability offers two predefined locations as well as a custom option that allows you to create your own path.
- **All To One**—This option replicates data from the source to a single volume on the target. The pre-defined path is \source_name\replication_set_name\volume_name. If you are replicating from multiple volumes on the source, each volume would be replicated to the same volume on the target. For example, c:\data and d:\files for the source Alpha and replication set DataFiles would be replicated to c:\alpha\DataFiles\c and c:\alpha\DataFiles\d, respectively.
- **One To One**—This option replicates data from the source to the same directory structure on the target. For example, c:\data and d:\files on the source will be replicated to c:\data and d:\files, respectively, on the target.
- **Custom Location**—If the predefined options do not store the data in a location that is appropriate for your network operations, you can specify your own custom location where the replicated files will be sent. Click the **Target Path** and edit it, selecting the appropriate location.

Note: If you are mirroring and replicating dynamic volumes or mount points, your location on the target must be able to accommodate the amount of data that you are replicating.

If you are mirroring and replicating sparse files and your location on the target is a non-NTFS 5 volume, the amount of disk space available must be equal to or greater than the entire size of the sparse file. If you are mirroring and replicating to an NTFS 5 volume, the

amount of disk space available must be equal to or greater than the on-disc size of the sparse file.

If you are mirroring and replicating multiple mount points, your directory mapping must not create a cycle or loop. For example, if you have the c: volume mounted at d:\c and the d: volume mounted at c:\d, this is a circular configuration. If you create and connect a replication set for either c:\d or d:\c, there will be a circular configuration and mirroring will never complete.

- **Start Mirror on Connection**—Mirroring can be initiated immediately when the connection is established. If mirroring is not configured to start automatically, you must start it manually after the connection is established.

Note: Data integrity cannot be guaranteed without a mirror being performed. This option is recommended for the initial connection.

- **Start Replication on Connection**—Replication can be initiated immediately when the connection is established. If replication is not configured to start automatically, you must start it manually after the connection is established. If you disable this option, you will need to perform a mirror prior to beginning replication to guarantee integrity.
- **Block Target Path(s) on Connection**—You can block writing to the data located in the target paths. This keeps the data from being changed outside of Double-Take Availability processing. If you are going to use failover, any target paths that are blocked will automatically be unblocked during the failover process so that users can modify data on the target after failover. During a restoration, the paths are automatically blocked again. If you failover and failback without performing a restoration, the target paths will remain unblocked. You can manually block or unblock the target paths by right-clicking on a connection.

3. If desired, you can configure mirror settings before establishing your connection. Select the **Mirroring** tab on the Connection Manager.

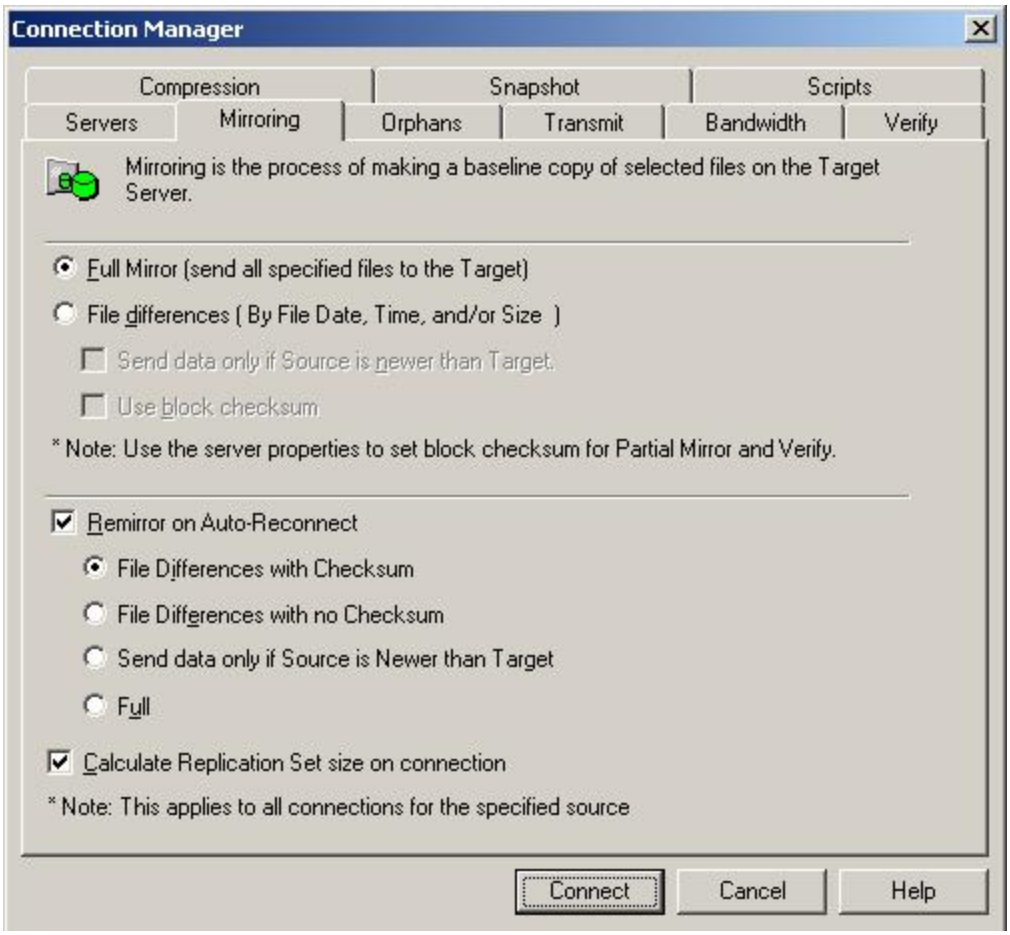

- **Full Mirror**—All files in the replication set will be sent from the source to the target.
- **File Differences**—Only those files that are different based on date, time, and/or size will be sent from the source to the target.
  - **Send data only if Source is newer than Target**—Only those files that are newer on the source are sent to the target.

Note: If you are using a database application, do not use the newer option unless you know for certain you need it. With database applications, it is critical that all files, not just some of them that might be newer, get mirrored.

  - **Use block checksum**—For those files flagged as different, the mirror performs a checksum comparison and only sends those blocks that are different.

Note: Stopping, starting, pausing, or resuming mirroring contains a comparison of how the file difference mirror settings work together, as well as how they work with the global checksum setting on the **Source** tab of the **Server Properties**.

- **Remirror on Auto-Reconnect**—In certain circumstances, for example if the disk-based queues on the source are exhausted, Double-Take Availability will automatically disconnect connections (called auto-disconnect) and then automatically reconnect them (called auto-reconnect). In order to ensure data integrity on the target, Double-Take Availability will perform an automatic mirror (called an auto-remirror) after an auto-reconnect. If you enable this option, specify the type of auto-remirror that will be performed.
  - **File Differences with Checksum**—Any file that is different on the source and target based on date, time, and/or size is flagged as different. The mirror then performs a checksum comparison on the flagged files and only sends those blocks that are different.
  - **File Differences with no Checksum**—Any file that is different on the source and target based on date, time, and/or size is sent to the target.
  - **Send data only if Source is newer than Target**—Only those files that are newer on the source are sent to the target.

Note: Stopping, starting, pausing, or resuming mirroring contains a comparison of how the file difference mirror settings work together, as well as how they work with the global checksum setting on the **Source** tab of the **Server Properties**.

  - **Full**—All files are sent to the target.

Note: Database applications may update files without changing the date, time, or file size. Therefore, if you are using database applications, you should use the **File Differences with checksum** or **Full** option.

- **Calculate Replication Set size on connection**—Determines the size of the replication set prior to starting the mirror. The mirroring status will update the percentage complete if the replication set size is calculated.

4. Click **Connect** to establish the connection.

Note: The settings on the other tabs of the Connection Manager are advanced settings. You can modify any of them before or after establishing your connection.

If you decide to enable orphan file processing while you are establishing your connection, orphan files will not be immediately processed when you create the connection. This setting is for processes that are run after a connection is already established (remirror, auto-remirror, verification, and so on).

Establishing a connection across a NAT or firewall

If your source and target are on opposite sides of a NAT or firewall, you will need special configurations to accommodate the complex network environment. Additionally, you must have the hardware already in place and know how to configure the hardware ports. If you do not, see the reference manual for your hardware.

In this environment, you must have static mapping where a single, internal IP address is always mapped in a one-to-one correlation to a single, external IP address. Double-Take Availability cannot handle dynamic mappings where a single, internal IP address can be mapped to any one of a group of external IP addresses managed by the router.

1. Double-Take Availability uses specific ports for communication between the Double-Take Availability servers and Double-Take Availability clients. In order to use Double-Take Availability through a NAT or firewall, you must first verify the current Double-Take Availability port settings so that you can open the correct ports on your hardware to allow Double-Take Availability machines to communicate with each other. By default, Double-Take Availability uses port 6320 for all communications. If you have changed your Double-Take Availability port, you will need to identify what port number is being used. The port setting can be found in the following locations.
    - **Replication Console**—From the Replication Console, select **File**, **Options**, and the **Configuration** tab.
    - **Failover Control Center**—From the Failover Control Center, select **Settings**, **Communications**.
    - **Double-Take Availability server**—From the Replication Console, right-click on a server in the tree in the left pane of the Replication Console, select **Properties**, and the **Network** tab.

    Note: If you change any of the port settings, you must stop and restart the Double-Take service for the new port setting to take effect.

    ---

2. You need to configure your hardware so that Double-Take Availability traffic is permitted access through the router and directed appropriately. Configure your router identifying each Double-Take Availability server, its IP address, and the Double-Take Availability and router ports. Also, note the following caveats.
    - Since Double-Take Availability communication occurs bidirectionally, make sure you configure your router for both incoming and outgoing traffic for all of your Double-Take Availability servers and Double-Take Availability clients.

- Double-Take Availability failover can use ICMP pings to determine if the source server is online. If you are going to use ICMP pings and a router between the source and target is blocking ICMP traffic, failover monitors cannot be created or used. In this situation, you must configure your router to allow ICMP pings between the source and target.

  Since there are many types of hardware on the market, each can be configured differently. See your hardware reference manual for instructions on setting up your particular router.

3. Manually insert the servers, by selecting **Insert**, **Server**. Type the IP address of the router the server is connected to and the port number the server is using for heartbeats.
4. Once your server is inserted in the Replication Console, use the Connection Manager to establish your connection. When specifying the Route on the Connection Manager **Servers** tab, you can manually enter the external IP address so traffic is routed appropriately.

Simulating a connection

Double-Take Availability offers a simple way for you to simulate a connection in order to generate statistics that can be used to approximate the time and amount of bandwidth that the connection will use when actively established. This connection uses the TDU (Throughput Diagnostics Utility), which is a built-in null (non-existent) target to simulate a real connection. No data is actually transmitted across the network. Since there is no true connection, this connection type helps you plan for a disaster recovery solution.

Before and after simulating your connection, you should gather network and system information specific to Double-Take Availability operations. Use the DTInfo utility to automatically collect this data. It gathers Double-Take Availability log files; Double-Take Availability and system settings; network configuration information such as IP, WINS and DNS addresses; and other data which may be necessary in evaluating Double-Take Availability performance. The DTInfo utility can be found on the product CD, in the Double-Take Availability installation directory, or on the Double-Take Software support web site.

1. From the source where you will be running the TDU, run DTInfo.exe. It may take several minutes for DTInfo to finish processing. After DTInfo processing is complete, a \support subdirectory will automatically be created in the Double-Take Availability installation directory. A .zip file will contain the information gathered. The file name is based on the machine name. To distinguish this file from the next time you run DTInfo, append a unique identifier, perhaps the date and time, to the end of the file name.
2. Make sure you have created a replication set that contains the data you want to protect.
3. From the Replication Console, open the Connection Manager to establish the connection.
   - Highlight the replication set and select **Tools**, **Connection Manager**.
   - Right-click on the replication set and select **Connection Manager**.
4. The Connection Manager opens to the **Servers** tab. Depending on how you opened the Connection Manager, some entries on the **Servers** tab will be completed already. For example, if you accessed the Connection Manager by right-clicking on a replication set, the name of the replication set will be displayed in the Connection Manager. Verify or complete the fields on the **Servers** tab.

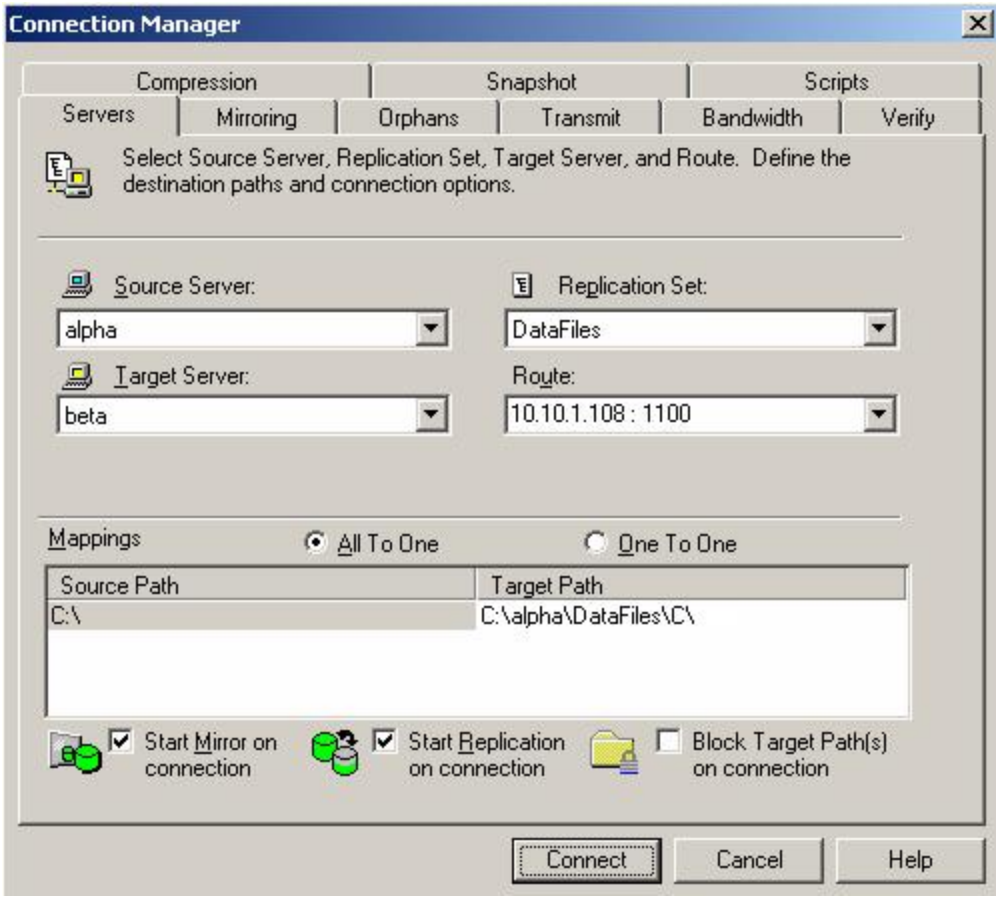

- **Source Server**—Specify the source server that contains the replication set that is going to be simulated to the TDU.
- **Replication Set**—At least one replication set must exist on the source before establishing a connection. Specify the replication set that will be connected to the TDU.
- **Target Server**—Select the **Diagnostics** target.
- **Route**—After selecting the **Diagnostics** target, the **Route** will automatically be populated with Throughput Diagnostics Utility (TDU).
- **Mappings**—Mappings are not required when simulating a connection because no data is actually transmitted to the target.
- **Start Mirror on Connection**—Make sure this option is selected so that your simulation will be realistic.
- **Start Replication on Connection**—Make sure this option is selected so that your simulation will be realistic.
- **Block Target Path(s) on Connection**—This option is not needed when simulating a connection because no data is actually transmitted to the target.

5. Click **Connect** to establish the connection. The simulation data will be logged to the Double-Take Availability statistics file.
6. Repeat step 1 to run the diagnostics utility after the simulation is complete.

Data workload failover

When you established your data workload protection, you only protected the data. You still need to establish failover monitoring, so that the target can stand in for the source in the event of a source failure, and your end users can access the data from the target.

- Configuring failover monitoring
- Updating shares on the target
- Editing failover monitoring configuration
- Removing failover monitoring configuration

Configuring failover monitoring

1. Open the Failover Control Center from your target server or a client machine. Do not run the Failover Control Center from your source server.
2. Select a failover target from the **Target Machine** list box.

Note: If the target you need is not listed, click **Add Target** and manually enter a name or IP address (with or without a port number). You can also select the **Browse** button to search for a target machine name. Click **OK** to select the target machine and return to the Failover Control Center main window.

---

3. Click **Login** to login to the selected target.
4. Select a source machine to monitor by clicking **Add Monitor**. The Insert Source Machine dialog box appears in front of the Monitor Settings dialog box.
5. On the Insert Source Machine dialog, specify your source machine by any of the following methods.

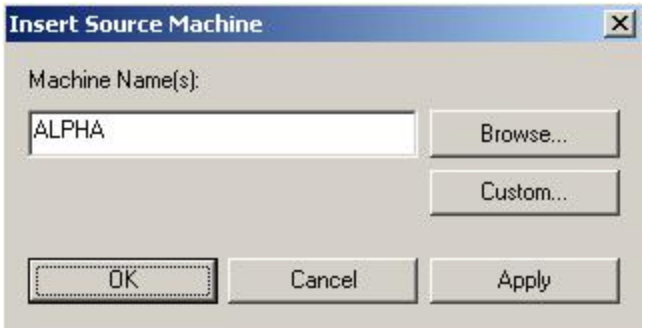

- Type the name of the machine that you want to monitor in Machine Name(s) and click **OK**.
- Click **Browse** to search for a machine. Select a domain from the list box at the top of the Select Machine dialog box to list the available machines for that domain. Highlight a source to be monitored and **click OK**.
- Click **Custom**. Enter the name of the server and click **Add**. Specify the IP address and subnet mask of the specified server and click **OK**. Click **OK** again.

  The Insert Source Machine dialog closes and the Monitor Settings dialog remains open with your source listed in the **Names to Monitor** tree.

6. In the **Names to Monitor** tree, locate and select the IP addresses on the source that you want to monitor.

7. Highlight an IP address that you have selected for monitoring and select a **Target Adapter** that will assume that IP address during failover. Repeat this process for each IP address that is being monitored. **Current IP Addresses** displays the IP address(es) currently assigned to the selected target adapter.
8. Highlight an IP address that you have selected for monitoring and select a **Method to Monitor for Failover**.
   - **Network Service**—Source availability will be tested for by a Double-Take Availability network response
   - **Replication Service**—Source availability will be tested for by a Double-Take service response.
   - **Network and Replication**—Source availability will be tested for by both a Double-Take Availability network response and a Double-Take service response.
   - **No Monitoring**—Double-Take Availability does not actively monitor the source. You will be responsible for identifying when a failure has occurred and initiating failover manually.
9. Repeat step 8 for each IP address that is being monitored.
10. Highlight an IP address that you have selected for monitoring and select a **Monitor Interval**. This setting identifies the number of seconds between the monitor requests sent from the target to the source to determine if the source is online. This option is not configurable if your **Method to Monitor for Failover** is set to **No Monitoring**. Repeat this step for each IP address that is being monitored.
11. Highlight an IP address that you have selected for monitoring and select the **Missed Packets**. This setting is the number of monitor replies sent from the source to the target that can be missed before assuming the source machine has failed. This option is not configurable if your **Method to Monitor for Failover** is set to **No Monitoring**. Repeat this step for each IP address that is being monitored.

Note: To achieve shorter delays before failover, use lower **Monitor Interval** and **Missed Packets** values. This may be necessary for servers, such as a web server or order processing database, which must remain available and responsive at all times. Lower values should be used where redundant interfaces and high-speed, reliable network links are available to prevent the false detection of failure. If the hardware does not support reliable communications, lower values can lead to premature failover. To achieve longer delays, choose higher values. This may be necessary for machines on slower networks or on a server that is not transaction critical. For example, failover would not be necessary in the case of a server restart.

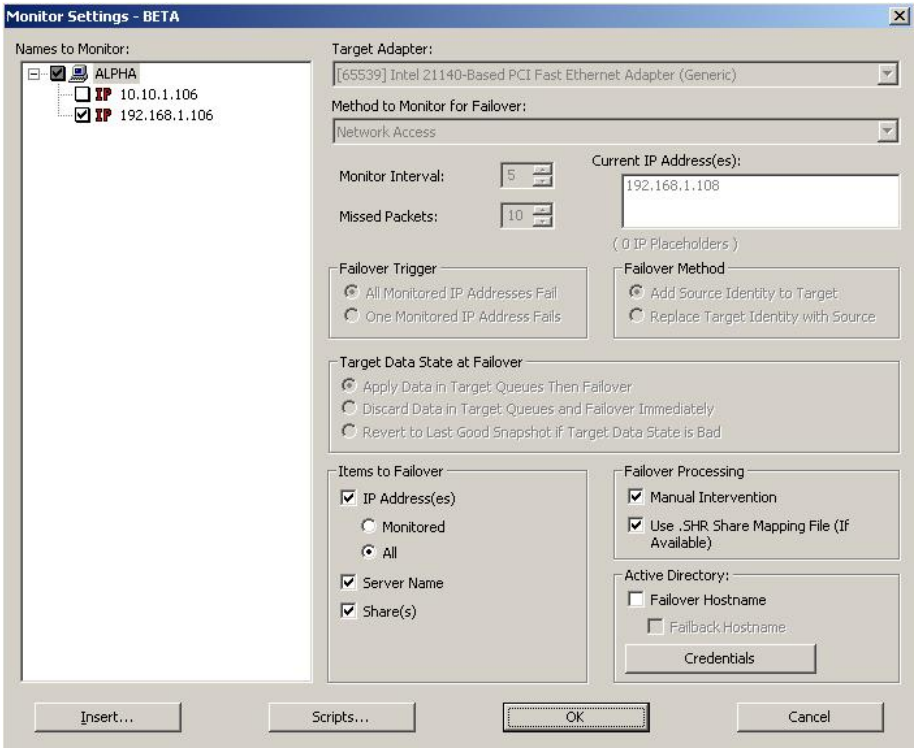

12. If you are monitoring multiple IP addresses, highlight the source name and specify the **Failover Trigger**.
    - **All Monitored IP Addresses Fail**—Failover begins when all monitored IP addresses fail. If there are multiple, redundant paths to a server, losing one probably means an isolated network problem and you should wait for all IP addresses to fail.
    - **One Monitored IP Address Fails**—Failover begins when any one of the monitored IP addresses fails. If each IP address is on a different subnet, you may want to trigger failover after one fails.
13. If **Manual Intervention** is enabled, **Target Data State at Failover** will be disabled because the same options will be presented to you at failover time. If **Manual Intervention** is disabled, the options will be enabled so that an option can be selected to occur automatically when failover occurs. Highlight the source name and specify the **Target Data State at Failover** by specifying what data you want to use on the target when failover occurs.
    - **Apply Data in Target Queues Then Failover**—All of the data in the target queue will be applied before failover begins. Depending on the amount of data in queue, the amount of time to apply all of the data could be lengthy.
    - **Discard Data in Target Queues and Failover Immediately**—All of the data in the target queue will be discarded and failover will begin immediately. Any data in the target queue will be lost.

- **Revert to Last Good Snapshot if Target Data is Bad**—If the target data is in a bad Double-Take Availability state, Double-Take Availability will automatically revert to the last good Double-Take Availability snapshot before failover begins. You will lose any data between the last good snapshot and the failure. If the target data is in a good state, Double-Take Availability will not revert the target data. Instead, Double-Take Availability will apply the data in the target queue and then failover. Depending on the amount of data in queue, the amount of time to apply all of the data could be lengthy.

14. Highlight the source name and specify the **Items to Failover**, which identifies which source components you want to failover to the target.
    - **IP Addresses**—If you want to failover the IP addresses on the source, enable this option and then specify the addresses that you want to failover.
    - **Monitored**—Only the IP address(es) that are selected for monitoring will be failed over.
    - **All**—All of the IP address(es) will be failed over.

    Note: If you are monitoring multiple IP addresses, IP address conflicts may occur during failover when the number of IP addresses that trigger failover is less than the number of IP addresses that are assumed by the target during failover. For example, if a source has four IP addresses (three public and one private), and two of the three public addresses are monitored, but all three public addresses are configured to failover, a conflict could occur. If the source fails, there is no conflict because all of the IP addresses have failed and no longer exist. But if the failure only occurs on one of the monitored addresses, the other two IP addresses are still affected. If all of the addresses are failed over, these addresses then exist on both the source and the target. Therefore, when a source machine has fewer IP addresses that trigger failover than IP addresses that will be failed over, there is a risk of an IP address conflict.

    If your network environment is a WAN configuration, do not failover your IP addresses unless you have a VPN infrastructure so that the source and target can be on the same subnet, in which case IP address failover will work the same as a LAN configuration. If you do not have a VPN, you can automatically reconfigure the routers via a failover script (by moving the source's subnet from the source's physical network to the target's physical network). There are a number of issues to consider when designing a solution that requires router configuration to achieve IP address failover. Since the route to the

source's subnet will be changed at failover, the source server must be the only system on that subnet, which in turn requires all server communications to pass through a router. Additionally, it may take several minutes or even hours for routing tables on other routers throughout the network to converge.

- **Server Name**—Failover is performed on the server name. If you specify the server name to be failed over, first Double-Take Availability checks the hosts file and uses the first name there. If there is no hosts file, Double-Take Availability uses the first name in DNS. (Although, the first name in DNS may not always be the same each time the DNS server is rebooted.) Lastly, if there is no DNS server, Double-Take Availability uses the Failover Control Center monitor name.
- **Shares**—Failover is performed on shares.

Note: Automatic share failover only occurs for standard Windows file system shares. Other shares must be configured for failover through the failover scripts or created manually on the target. See Macintosh shares or NFS Shares for more information.

If you are failing over Windows shares but your source and target do not have the same drive letters, you must use the **All to One** selection when establishing your Double-Take Availability connection. Otherwise, the shares will not be created on the target during failover.

If a Windows share is created on Windows 2003 with the default full access permissions (without an ACL) and then failed over, the permissions given to the target will be read-only permissions.

Windows share information is automatically updated on the target once an hour. If you need to manually update the share information, click Update Shares on the main Failover Control Center window after the monitor has been established.

---

15. By default, **Manual Intervention** is enabled, allowing you to control when failover occurs. When a failure occurs, a prompt appears in the Failover Control Center and waits for you to manually initiate the failover process. Disable this option only if you want failover to occur immediately when a failure occurs. This option is not configurable if the **Method to Monitor for Failover** is set to **No Monitoring**.
16. If the **Shares** selection under **Items to Failover** is selected, verify that the **Use**

.SHR Share Mapping File check box is selected if you would like to use the Double-Take Availability share mapping file to create shares on the target during failover. If this option is not selected, shares will be created using the information gathered when the machine was selected as a source to be monitored.

Note: If the **Shares** selection under **Items to Failover** is not selected, shares will not be failed over to the target regardless of the **Use .SHR Share Mapping File** selection.

---

17. By default, **Failover Hostname** is disabled. This option automatically removes the host SPN (Service Principle Name) from Active Directory on the source and adds it to Active Directory on the target. If you are using Active Directory, enable this option or you may experience problems with failover.
18. **Failback Hostname** returns the host SPN on the source and target back to their original settings on failback. If you are using Active Directory, enable this option or you may experience problems with failback.
19. If you are failing over or failing back hostnames, you need to specify an Active Directory user that has update privileges within Active Directory. Click **Credentials** and identify a user and the associated password that has privileges to create and delete SPNs. The username must be in the format fully_qualified_domain\user. Click **OK** to return to the Monitor Settings dialog box.

Note: The Active Directory account password cannot be blank.

---

20. If you are using any failover or failback scripts, click **Scripts** and enter the path and filename for each script type. Scripts may contain any valid Windows command, executable, or batch file. Examples of functions specified in scripts include stopping services on the target before failover because they may not be necessary while the target is standing in for the source, stopping services on the target that need to be restarted with the source's machine name and IP address, starting services or loading applications that are in an idle, standby mode waiting for failover to occur, notifying the administrator before and after failover or failback occurs, stopping services on the target after failback because they are no longer needed, stopping services on the target that need to be restarted with the target machine's original name and IP address, and so on. Specify each script that you want to run and the following options, if necessary.

- If you want to pass any arguments to your script, specify the arguments.
- If you want to delay the failover or failback processes until the associated script has completed, mark the appropriate check box.
- If you want the same scripts to be used as the default for future monitor sessions, mark the appropriate check box.
- If you want to specify a user account to run the scripts, specify the credentials for the source and target. If no account is specified, the scripts will be processed using the same account that is configured to run the Double-Take service.

21. Click **OK** to return to the Monitor Settings dialog box.

Note: Failover scripts will run but will not be displayed on the screen if the Double-Take service is not set to interact with the desktop. Enable this option through the Services applet.

With these flexible scripting features, application failover using Double-Take Availability can be seamless to the end user. Double-Take Software tests many of the popular applications on the market today. The results of these testing procedures are written up into formal Application Notes that describe how Double-Take Availability should be configured to work correctly with certain applications, including scripting examples. For a complete list of Double-Take Availability Application Notes, visit the Double-Take Software support web site.

---

22. Click **OK** on the Monitor Settings dialog box to save your monitor settings and begin monitoring for a failure.

Updating shares on the target

Share information on the target can be manually updated from the Failover Control Center window. To manually update the share information, highlight a source machine in the **Monitored Machines** tree and click the **Update Shares** button.

Editing failover monitoring configuration

If you want to edit the monitor settings for a source that is currently being monitored, highlight that source on the **Monitored Machines** tree on the main Failover Control Center screen and click **Edit**. The Monitor Settings dialog box will open. Follow the Configuring failover monitoring instructions.

Removing failover monitoring configuration

If you want to discontinue monitoring a source, highlight that machine on the **Monitored Machines** tree on the main Failover Control Center screen and click **Remove Monitor**. No additional dialog boxes will open.

Server settings

Most of the Double-Take Availability server settings are located in the Replication Console Server Properties dialog box. To access this dialog box, right-click a server in the left pane of the Management Console and select **Properties**. The Server Properties dialog box contains multiple tabs with the Double-Take Availability server settings. For information on the server settings not available through the Replication Console, see the *Scripting Guide*.

This section contains the following topics, each corresponding to a tab in the Server Properties dialog box.

- Identifying a server
- Licensing a server
- Configuring server startup options
- Configuring network communication properties for a server
- Queuing data
- Configuring source data processing options
- Configuring target data processing options
- Specifying the Double-Take Availability database storage files
- Specifying file names for logging and statistics
- Supplying credentials for script processing
- E-mailing event messages

Identifying a server

From the Replication Console, you can see server identity information, including a server's Double-Take Availability activation code.

1. Open the Replication Console and right-click the server on the left pane of the Replication Console.
2. Select **Properties**
3. Select the **General** tab.

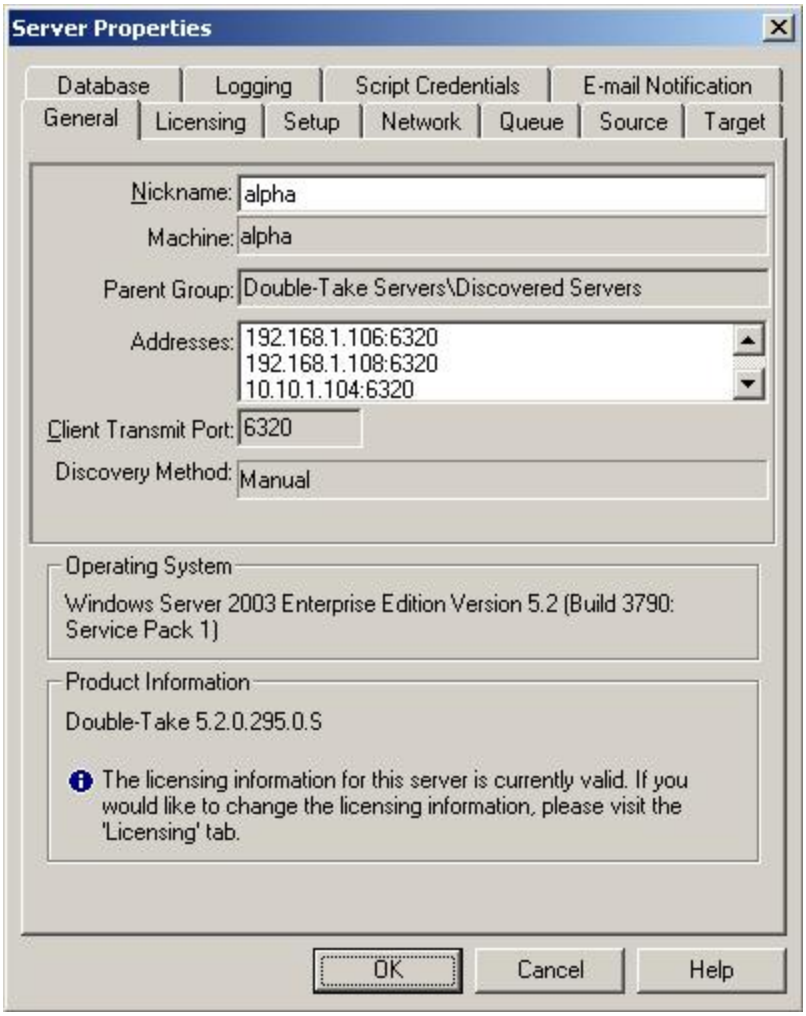

4. Specify the server identity information. Some of the fields are informational only.
   - **Nickname**—A nickname is saved in the Replication Console workspace, therefore, it only appears in the Replication Console on this server. It is not communicated across the network. If you export a workspace and use it on another Double-Take Availability server, the server nickname will appear there also.
   - **Machine**—This is the actual server name. This field is not modifiable.

- **Parent Group**—This is the group in the Replication Console server
- **Addresses**—The IP address(es) for this server are listed in this field. This information is not modifiable and is displayed for your information. The machine's primary address is listed first.
- **Client Transmit Port**—This field displays the port that the Replication Console uses to send commands to a server. This port cannot be modified.
- **Discovery Method**—This field indicates the method in which the Replication Console identifies the Double-Take Availability server. **Manual** indicates a server was manually inserted into the Replication Console server tree. **Active Directory** indicates a server is registered with Windows Active Directory.
- **Operating System**—The server's operating system version is displayed.
- **Product Information**—The Double-Take service name and the build number are displayed.

5. Click **OK** to save the settings.

Licensing a server

From the Replication Console, you can manage your server activation codes. The activation code is the Double-Take Availability license which is required on every Double-Take Availability server. The activation code is a 24 character, alpha-numeric code. You can change your activation code without reinstalling, if your license changes. There are different licenses available.

- **Evaluation**—An evaluation license has an expiration date built into the activation code. When the license expires, the software will no longer function. The same evaluation licenses can be used on multiple machines on a network.
- **Single**—A single license is available on a per-machine basis. Each server is required to have a unique license whether it is functioning as a source, target, or both. A single license can only be used on one server on a network.
- **Site**—A site license is available to register every machine with the same license. This license is designed to be used on multiple servers on a network.
- **Node-Locking**—To prevent Double-Take Availability from being used illegally on multiple servers, you may have received a node-locked activation code, which is a temporary license. The temporary license is not activated until you login to the server. Once the temporary license is activated, you have 14 days to update it to a permanent, node-locked license. The permanent node-locked license will be created by supplying unique server information to Double-Take Software. Since the permanent node-locked license contains unique server information, specific to the hardware where Double-Take Availability is installed, the node-locked license cannot be used on any other server, thus prohibiting illegal applications.

Use the following instructions to access licensing information.

1. Open the Replication Console and right-click the server on the left pane of the Replication Console.
2. Select **Properties**.
3. Select the **Licensing** tab. The fields displayed on this tab will vary depending on

your activation code(s).

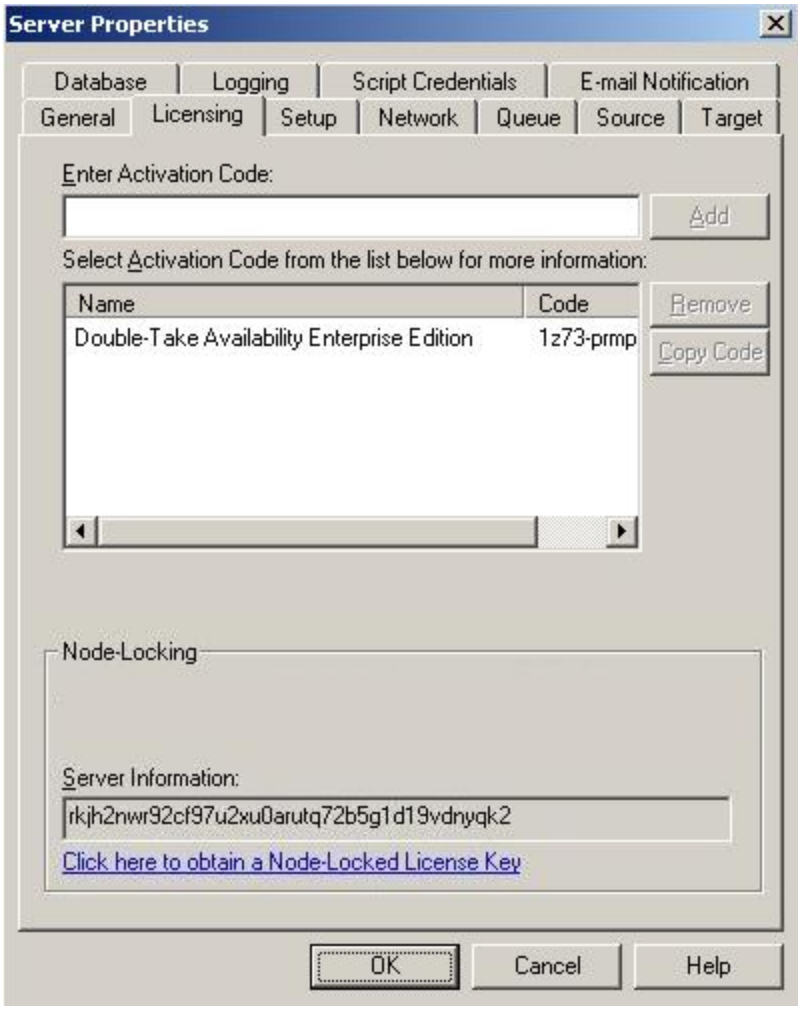

4. Enter an activation code and click **Add**. Repeat for each activation code.
5. Highlight an activation code in the list to display any status messages for that code below the list display.
6. If you need to remove a code from the server, highlight it in the list and click **Remove**.
7. To update a temporary node-locked license to a permanent license, you need to provide server information which will be used to generate a permanent node-locked license.
8. Click **OK** to save the settings.

Configuring server startup options

From the Replication Console, you can configure server startup options for each Double-Take Availability server.

1. Open the Replication Console and right-click the server on the left pane of the Replication Console.
2. Select **Properties**
3. Select the **Setup** tab.

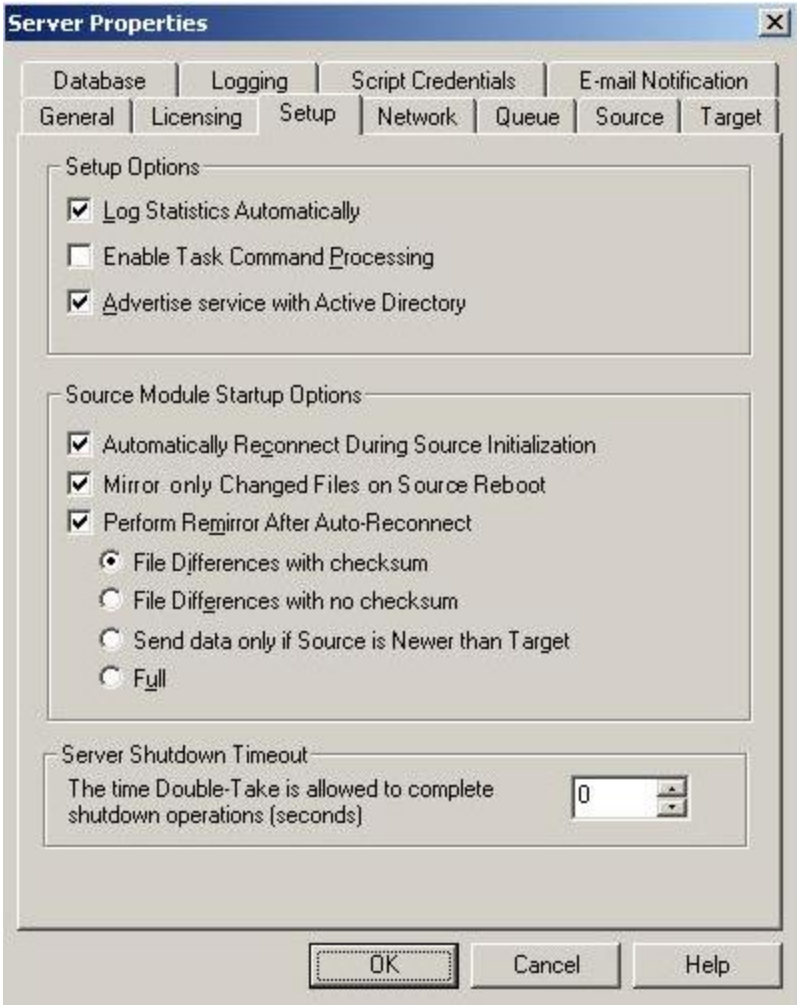

4. Specify the server setup and source startup options.
   - **Log Statistics Automatically**—If enabled, Double-Take Availability statistics logging will start automatically when Double-Take Availability is started.
   - **Enable Task Command Processing**—Task command processing is a Double-Take Availability feature that allows you to insert and run tasks at

various points during the replication of data. Because the tasks are user-defined, you can achieve a wide variety of goals with this feature. For example, you might insert a task to create a snapshot or run a backup on the target after a certain segment of data from the source has been applied on the target. This allows you to coordinate a point-in-time backup with real-time replication.

Task command processing can be enabled from the Replication Console, but it can only be initiated through the scripting language. See the *Scripting Guide* for more information.

If you disable this option on a source server, you can still submit tasks to be processed on a target, although task command processing must be enabled on the target.

- **Advertise service with Active Directory**—If enabled, the Double-Take service registers with Windows Active Directory when the service is started.
- **Automatically Reconnect During Source Initialization**—If enabled, Double-Take Availability will automatically reconnect any connections that it automatically disconnected.
- **Mirror only Changed Files on Source Reboot**—If enabled, Double-Take Availability will use the Windows NTFS change journal to track file changes. If the source is rebooted, only the files identified in the change journal will be remirrored to the target. This setting helps improve mirror times.
- **Perform Remirror After Auto-reconnect**—If enabled, Double-Take Availability will automatically perform a remirror after an auto-reconnect has occurred. You will also need to specify the type of mirror that you wish to perform after an auto-reconnect.
  - **File Differences with Checksum**—Any file that is different on the source and target based on date, time, and/or size is flagged as different. The mirror then performs a checksum comparison on the flagged files and only sends those blocks that are different.
  - **File Differences with no Checksum**—Any file that is different on the source and target based on date, time, and/or size is sent to the target.
  - **Send data only if Source is newer than Target**—Only those files that are newer on the source are sent to the target.
  - **Full**—All files are sent to the target.

Note: Database applications may update files without changing the date, time, or file size. Therefore, if you are using database

applications, you should use the **File Differences with checksum** or **Full** option.

- **Server Shutdown Timeout**—This setting indicates the amount of time, in seconds, for the service to wait prior to completing a shutdown so that Double-Take Availability can persist data on the target in an attempt to avoid a remirror when the target comes back online. A timeout of zero (0) indicates waiting indefinitely and any other number indicates the number of seconds. The timeout setting only controls the service shutdown from the Double-Take Availability clients. It does not control the service shutdown through a reboot or from the Service Control Manager.

5. Click **OK** to save the settings.

Configuring network communication properties for a server

1. Open the Replication Console and right-click the server on the left pane of the Replication Console.
2. Select **Properties**
3. Select the **Network** tab.

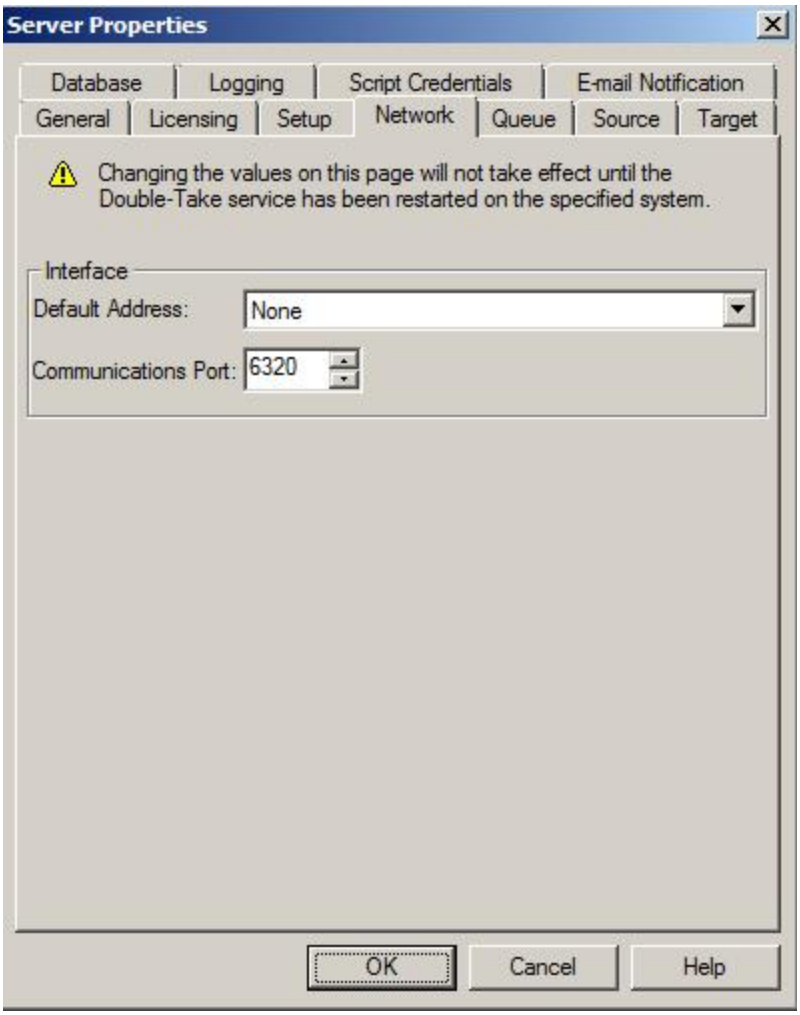

4. Specify the network communication properties.
   - **Default Address**—On a machine with multiple NICs, you can specify which address Double-Take Availability traffic will use. It can also be used on machines with multiple IP addresses on a single NIC.
   - **Communications port**—Double-Take Availability servers use this port to send and receive commands and operations between two Double-Take Availability servers.
5. Click **OK** to save the settings.

Queuing data

You should configure queuing on both the source and target.

1. Open the Replication Console and right-click the server on the left pane of the Replication Console.
2. Select **Properties**.
3. Select the **Queue** tab.
4. Specify the queue settings for the server.

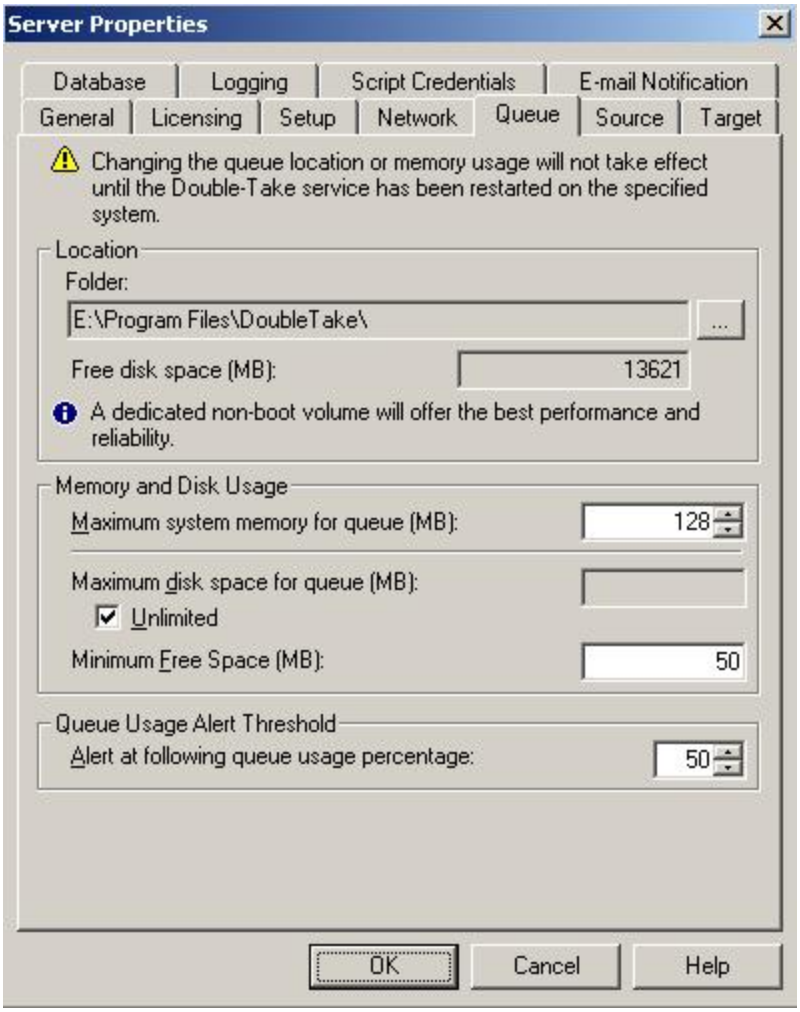

- **Folder**—This is the location where the disk queue will be stored. Double-Take Availability displays the amount of free space on the volume selected. Any changes made to the queue location will not take effect until the Double-Take service has been restarted on the server.

  When selecting the queue location, keep in mind the following caveats.

- Select a location on a non-clustered volume that will have minimal impact on the operating system and applications being protected.
- Select a location that is on a different volume as the location of the Windows pagefile.
- Select a dedicated, non-boot volume.
- Do not select the root of a volume.
- Do not select the same physical or logical volume as the data being replicated.

Although the read/write ratio on queue files will be 1:1, optimizing the disk for write activity will benefit performance because the writes will typically be occurring when the server is under a high load, and more reads will be occurring after the load is reduced. Accordingly, use a standalone disk, mirrored (RAID 1) or non-parity striped (RAID 0) RAID set, and allocate more I/O adapter cache memory to writes for best performance. A RAID 5 array will not perform as well as a mirrored or non-parity striped set because writing to a RAID 5 array incurs the overhead of generating and writing parity data. RAID 5 write performance can be up to 50% less than the write performance of a single disk, depending on the adapter and disk.

Another option is to use a solid state disk, which are hard drives that use RAM instead of disk platters. These devices are typically quite costly, but they will provide superior performance as a queuing device when the best performance is required.

Note: Scanning the Double-Take Availability queue files for viruses can cause unexpected results. If anti-virus software detects a virus in a queue file and deletes or moves it, data integrity on the target cannot be guaranteed. As long as you have your anti-virus software configured to protect the actual production data, the anti-virus software can clean, delete, or move an infected file and the clean, delete, or move will be replicated to the target. This will keep the target from becoming infected and will not impact the Double-Take Availability queues.

- **Maximum system memory for queue**—This is the amount of Windows system memory, in MB, that will be used to store data in queues. When exceeded, queuing to disk will be triggered. This value is dependent on the amount of physical memory available but has a minimum of 32 MB. By default, 128 MB or 512 MB of memory is used, depending on your operating

system. If you set it lower, Double-Take Availability will use less system memory, but you will queue to disk sooner which may impact system performance. If you set it higher, Double-Take Availability will maximize system performance by not queuing to disk as soon, but the system may have to swap the memory to disk if the system memory is not available.

Since the source is typically running a production application, it is important that the amount of memory Double-Take Availability and the other applications use does not exceed the amount of RAM in the system. If the applications are configured to use more memory than there is RAM, the system will begin to swap pages of memory to disk and the system performance will degrade. For example, by default an application may be configured to use all of the available system memory when needed, and this may happen during high-load operations. These high-load operations cause Double-Take Availability to need memory to queue the data being changed by the application. In this case, you would need to configure the applications so that they collectively do not exceed the amount of RAM on the server. Perhaps on a server with 1 GB of RAM running the application and Double-Take Availability, you might configure the application to use 512 MB and Double-Take Availability to use 256 MB, leaving 256 MB for the operating system and other applications on the system. Many server applications default to using all available system memory, so it is important to check and configure applications appropriately, particularly on high-capacity servers.

Any changes to the memory usage will not take effect until the Double-Take service has been restarted on the server.

- **Maximum disk space for queue**—This is the maximum amount of disk space, in MB, in the specified **Folder** that can be used for Double-Take Availability disk queuing, or you can select **Unlimited** which will allow the queue usage to automatically expand whenever the available disk space expands. When the disk space limit is reached, Double-Take Availability will automatically begin the auto-disconnect process. By default, Double-Take Availability will use an unlimited amount of disk space. Setting this value to zero (0) disables disk queuing.
- **Minimum Free Space**—This is the minimum amount of disk space in the specified **Folder** that must be available at all times. By default, 50 MB of disk space will always remain free. The **Minimum Free Space** should be less than the amount of physical disk space minus **Maximum disk space for queue**.

Note: The **Maximum disk space for queue** and **Minimum Free Space** settings work in conjunction with each other. For example, assume your queues are stored on a 10 GB disk with the **Maximum disk space** for queue set to 10 GB and the **Minimum Free Space** set to 500 MB. If another program uses 5 GB, Double-Take Availability will only be able to use 4.5 GB so that 500 MB remains free.

- **Alert at following queue usage percentage**—This is the percentage of the disk queue that must be in use to trigger an alert message in the Linux system log. By default, the alert will be generated when the queue reaches 50%.

5. Click **OK** to save the settings.

Configuring source data processing options

1. Open the Replication Console and right-click the server on the left pane of the Replication Console.
2. Select **Properties**
3. Select the **Source** tab.

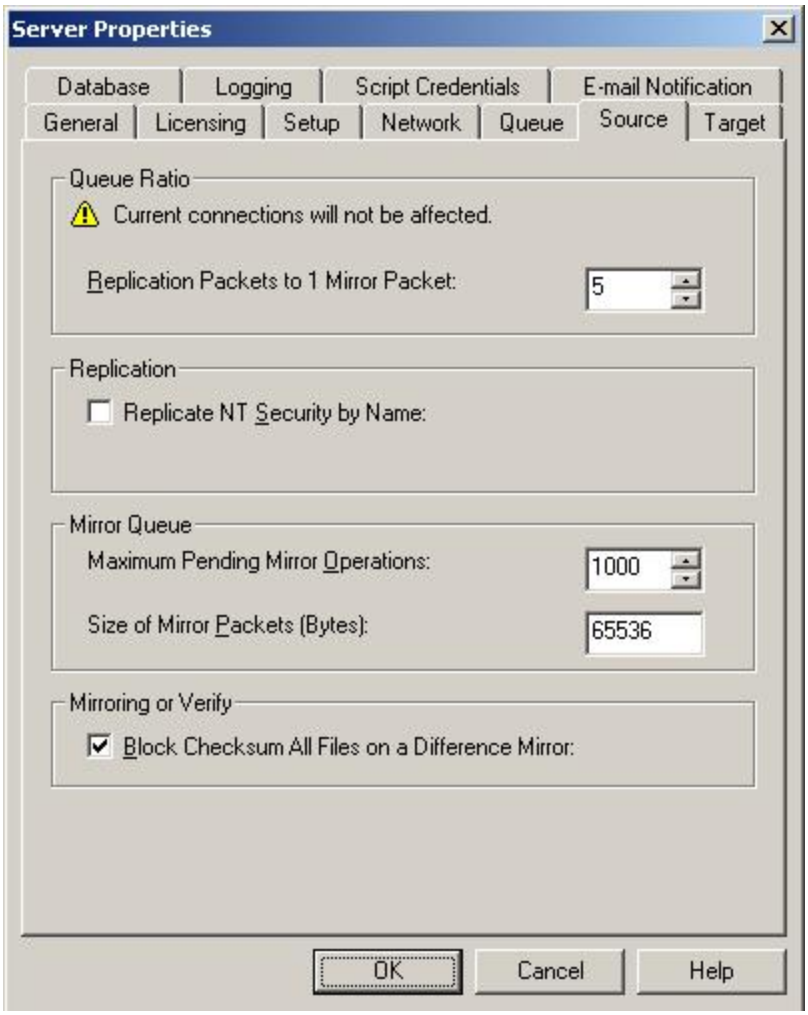

4. Specify how the source will process data.
   - **Replication Packets to 1 Mirror Packet**—You can specify the ratio of replication packets to mirror packets that are placed in the source queue. Specify a larger number if you have a busy network that has heavy replication. Also, if you anticipate increased network activity during a mirror, increase this number so that the replication queue does not get too large.
   - **Replicate NT Security by Name**—Double-Take Availability allows you to replicate Windows permission attributes by local name as well as security ID (SID). By replicating Windows security by name, you can transmit the owner

name with the file. If that user exists on the target, then the SID associated with the user will be applied to the target file ownership. If that user does not exist on the target, then the ownership will be unknown. By default, this option is disabled.

- **Domain security model**—If you are using a Windows domain security model by assigning users at the domain level, each user is assigned a security ID (SID) at the domain level. When Double-Take Availability replicates a file to the target, the SID is also replicated. Since a user will have the same SID on the source and target, the user will be able to access the file from the target. Therefore, this option is not necessary.
- **Local security model**—If you are using a Windows local security model by assigning users at the local level (users that appear on multiple machine will each have different SIDs), you will need to enable this feature so that users can access the data on the target. If you do not enable this feature with a local security model, after a Double-Take Availability file and SID is replicated, a local user will not be able to access the file because the user's SID on the target machine is different from the SID that was replicated from the source machine.

  If you enable this option, make sure that the same groups and users exist on the target as they do on the source. Additionally, you must enable this option on your target server before starting a restoration, because the target is acting like a source during a restoration.

  Enabling this option may have an impact on the rate at which Double-Take Availability can commit data on the target. File security attributes are sent to the target during mirroring and replication. The target must obtain the security ID (SID) for the users and groups that are assigned permissions, which takes some time. If the users and groups are not on the target server, the delay can be substantial. The performance impact enabling this option will have will vary depending on the type of file activity and other variables. For instance, it will not affect the overall performance of large database files much (since there is a lot of data, but only a few file permissions), but may affect the performance of user files significantly (since there are often thousands of files, each with permissions). In general, the performance impact will only be noticed during mirrors since that is when the target workload is greatest.

  Regardless of the security model you are using, if you create new user accounts on the source, you should start a remirror so the new user account information associated with any files in your replication set can be transmitted to the target.
- **Maximum Pending Mirror Operations**—This option is the maximum number of mirror operations that are queued on the source. The default

setting is 1000. If, during mirroring, the mirror queued statistic regularly shows low numbers, for example, less than 50, this value can be increased to allow Double-Take Availability to queue more data for transfer.

- **Size of Mirror Packets**—This option determines the size of the mirror packets that Double-Take Availability transmits. The default setting is 65536 bytes. You may want to consider increasing this value in a high latency environment (greater than 100 ms response times).
- **Block Checksum All Files on a Difference Mirror**—This option allows a file difference mirror to check each block of data, regardless of the file attributes. If this option is not marked, Double-Take Availability will assume files are synchronized if their attributes match.

Note: Database applications may update files without changing the date, time, or file size. Therefore, if you are using database applications, you should use the **Block Checksum All Files on a Difference Mirror** option to ensure proper file comparisons.

If you are not using database applications, disabling this option will shorten mirror times.

---

5. Click **OK** to save the settings.

Configuring target data processing options

1. Open the Replication Console and right-click the server on the left pane of the Replication Console.
2. Select **Properties**
3. Select the **Target** tab.

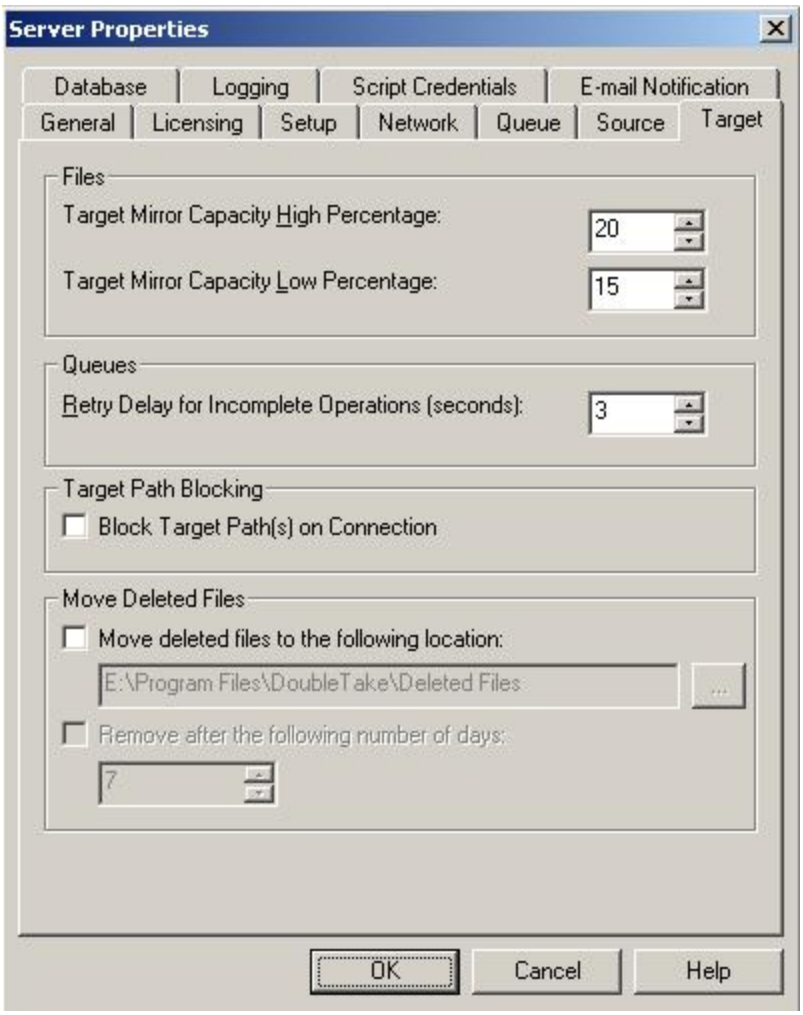

4. Specify how the target will process data.
   - **Target Mirror Capacity High Percentage**—You can specify the maximum percentage of Windows system memory that can contain mirror data before the target signals the source to pause the sending of mirror operations. The default setting is 20.
   - **Target Mirror Capacity Low Percentage**—You can specify the minimum percentage of Windows system memory that can contain mirror data before the target signals the source to resume the sending of mirror operations. The default setting is 15.

- **Retry Delay for Incomplete Operations (seconds)**—This option specifies the amount of time, in seconds, before retrying a failed operation on the target. The default setting is 3.
- **Block Target Path(s) on Connection**—You can block writing to the data located in the target paths. This keeps the data from being changed outside of Double-Take Availability processing.

Note: If you are going to use failover, any target paths that are blocked will automatically be unblocked during the failover process so that users can modify data on the target after failover. During a restoration, the paths are automatically blocked again. If you failover and failback without performing a restoration, the target paths will remain unblocked. You can manually block or unblock the target paths by right-clicking on a connection.

Do not block your target paths if you are protecting a full-server workload because system state data will not be able to be written to the target.

- **Move deleted files to the following location**—This option allows you to save files that have been deleted, by moving them to a different location on the target. When a file deletion is replicated to the target, instead of the file being deleted from the target, the file is moved to the specified location. This allows for easy recovery of those files, if needed. If you enable this option, specify where you want to store the deleted files.
- **Remove after the following number of days**—If you are moving deleted files, you can specify a length of time, in days, to maintain the moved files. A moved file that is older than the specified number of days will be deleted. Double-Take Availability checks for moved files that should be deleted once daily at 8 PM. Only the date, not the time, of the file is considered when moved files are deleted. For example, if you specify to delete moved files after 30 days, any file that is 31 days old will be deleted. Because the criteria is based on days and not time, a file that will be deleted could have been moved anytime between 12:01 AM and 11:59 PM 31 days ago.

Note: If deleted files are moved for long enough, the potential exists for the target to run out of space. In that case, you can manually delete files from the target move location to free space.

If you have **Ignore Delete Operations** enabled on the source, the **Move deleted files** setting will never be invoked because there will be no delete operations performed on the target.

Do not include the Recycler directory in your replication set if you are moving deleted files. If the Recycler directory is included, Double-Take Availability will see an incoming file deletion as a move operation to the Recycle Bin and the file will not be moved as indicated in the **Move deleted files** setting.

Alternate data streams that are deleted on the source will not be moved on the target.

Encrypted files that are deleted on the source will only be moved on the target if the move location is on the same volume as the replication set target path.

Compressed and sparse files that are deleted on the source will be moved on the target, although the compression and sparse flags will only be retained on the target if the move location is on the same volume as the replication set target path.

5. Click **OK** to save the settings.

Specifying the Double-Take Availability database storage files

1. Open the Replication Console and right-click the server on the left pane of the Replication Console.
2. Select **Properties**
3. Select the **Database** tab.

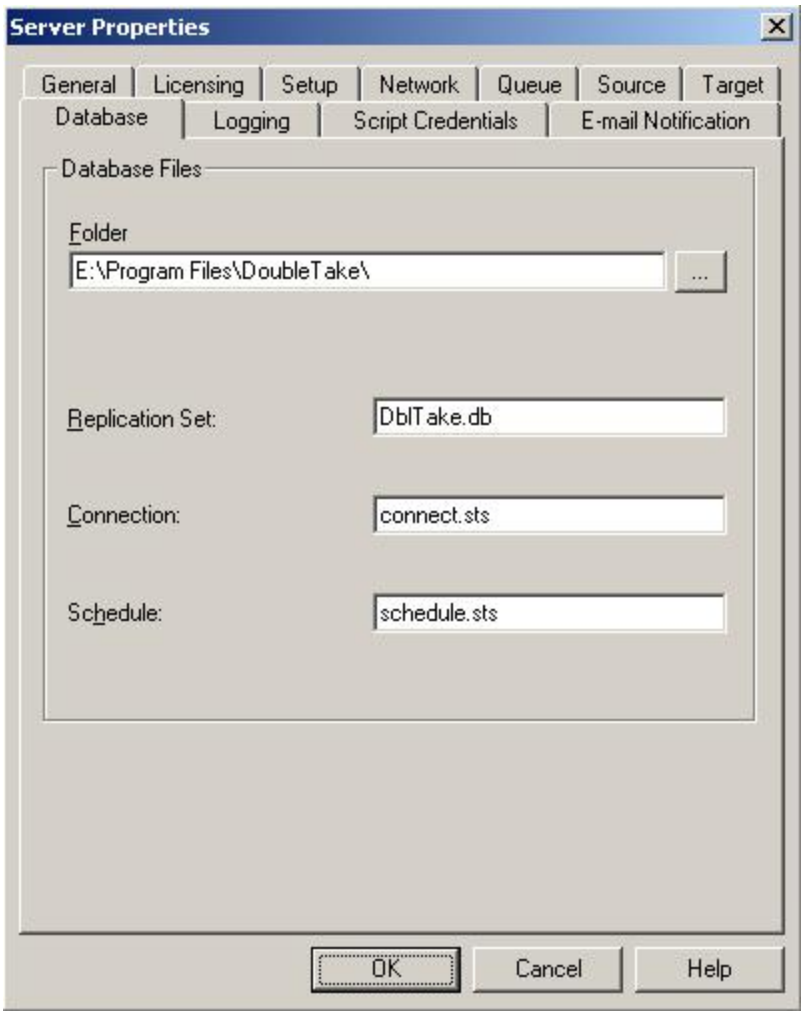

4. Specify the database files that store the Double-Take Availability replication set, connection, and scheduling information.
   - **Folder**—Specify the directory where each of the database files on this tab are stored. The default location is the directory where the Double-Take Availability program files are installed.
   - **Replication Set**—This database file maintains which replication sets have been created on the server along with their names, rules, and so on. The default file name is DblTake.db.

- **Connection**—This database file maintains the active source/target connection information. The default file name is connect.sts.
- **Schedule**—This database file maintains any scheduling and transmission limiting options. The default file name is schedule.sts.

5. Click **OK** to save the settings.

Specifying file names for logging and statistics

1. Open the Replication Console and right-click the server on the left pane of the Replication Console.
2. Select **Properties**
3. Select the **Logging** tab.

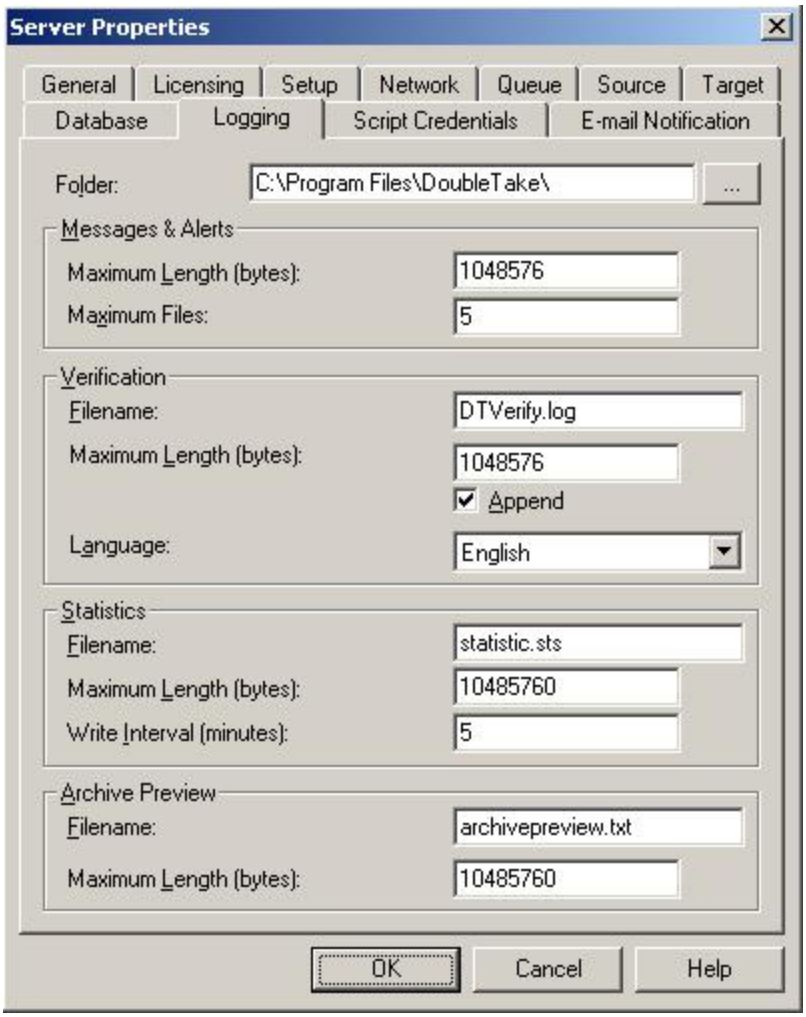

4. Specify the location and file names for the log and statistics files.
   - **Folder**—Specify the directory where each of the log files on this tab are stored. The default location is the directory where the Double-Take Availability program files are installed.
   - **Messages & Alerts**, **Maximum Length**—Specify the maximum length of the client and service log files. The default size is 1048576 bytes and is limited by the available hard drive space.

- **Messages & Alerts**, **Maximum Files**—Specify the maximum number of Double-Take Availability alert log files that are maintained. The default is 5, and the maximum is 999.
- **Verification**, **Filename**—The verification log is created during the verification process and details which files were verified as well as the files that are synchronized. This field contains the base log file name for the verification process. The replication set name will be prefixed to the base log file name. For example, since the default is DTVerify.log, the verification log for the replication set called UserData would be UserData DTVerify.log.
- **Verification**, **Maximum Length**—Specify the maximum length of the verification log file. The default maximum length is 1048576 bytes (1 MB).
- **Verification**, **Append**—Mark the Append check box if you want to append each verification process to the same log file. If this check box is not marked, each verification process that is logged will overwrite the previous log file. By default, this check box is selected.
- **Verification**, **Language**—At this time, English is the only language available.
- **Statistics**, **Filename**—The statistics log maintains connection statistics such as mirror bytes in queue or replication bytes sent. The default file name is statistic.sts. This file is a binary file that is read by the DTStat utility.
- **Statistics**, **Maximum Length**—Specify the maximum length of the statistics log file. The default maximum length is 10485760 bytes (10 MB). Once this maximum has been reached, Double-Take Availability begins overwriting the oldest data in the file.
- **Statistics**, **Write Interval**—Specify how often Double-Take Availability writes to the statistics log file. The default is every 5 minutes.
- **Archive Preview**, **Filename**—If you are using Double-Take Backup product, the archive preview identifies what files would be archived based on the specified criteria. The default file name is ArchivePreview.txt. The replication set name is appended to this filename when the report is generated.
- **Archive Preview**, **Maximum Length**—Specify the maximum length of the archive preview report. The default maximum length is 10485760 bytes (10 MB). Once this maximum has been reached, Double-Take Availability begins overwriting the oldest data in the file.

5. Click **OK** to save the settings.

Supplying credentials for script processing

1. Open the Replication Console and right-click the server on the left pane of the Replication Console.
2. Select **Properties**
3. Select the **Script Credentials** tab.

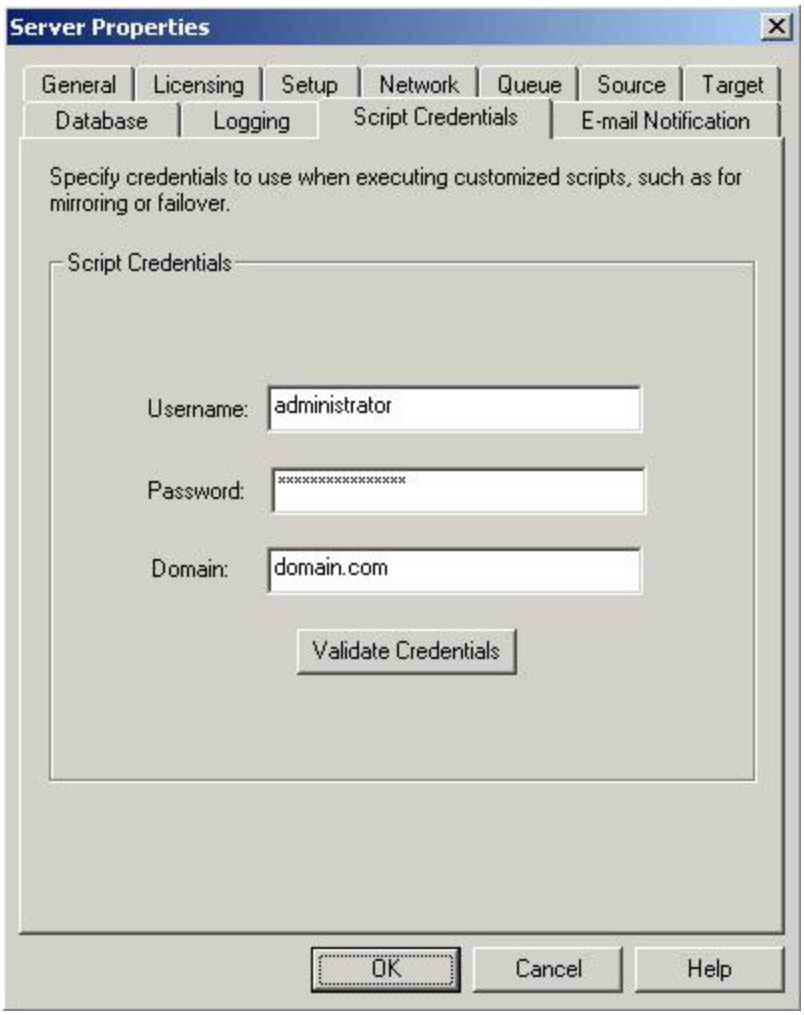

4. If you will be using any customized scripts (mirroring, failover, task command processing, and so on) specify a **Username**, **Password**, and **Domain** to use when running the scripts. If you do not specify any security credentials, the account running the Double-Take service will be used.
5. Click **OK** to save the settings.

E-mailing event messages

You can e-mail Double-Take Availability event messages to specific addresses. The subject of the e-mail will contain an optional prefix, the server name where the message was logged, the message ID, and the severity level (information, warning, or error). The text of the message will be displayed in the body of the e-mail message.

1. Open the Replication Console and right-click the server on the left pane of the Replication Console.
2. Select **Properties**.
3. Select the **E-mail Notification** tab.

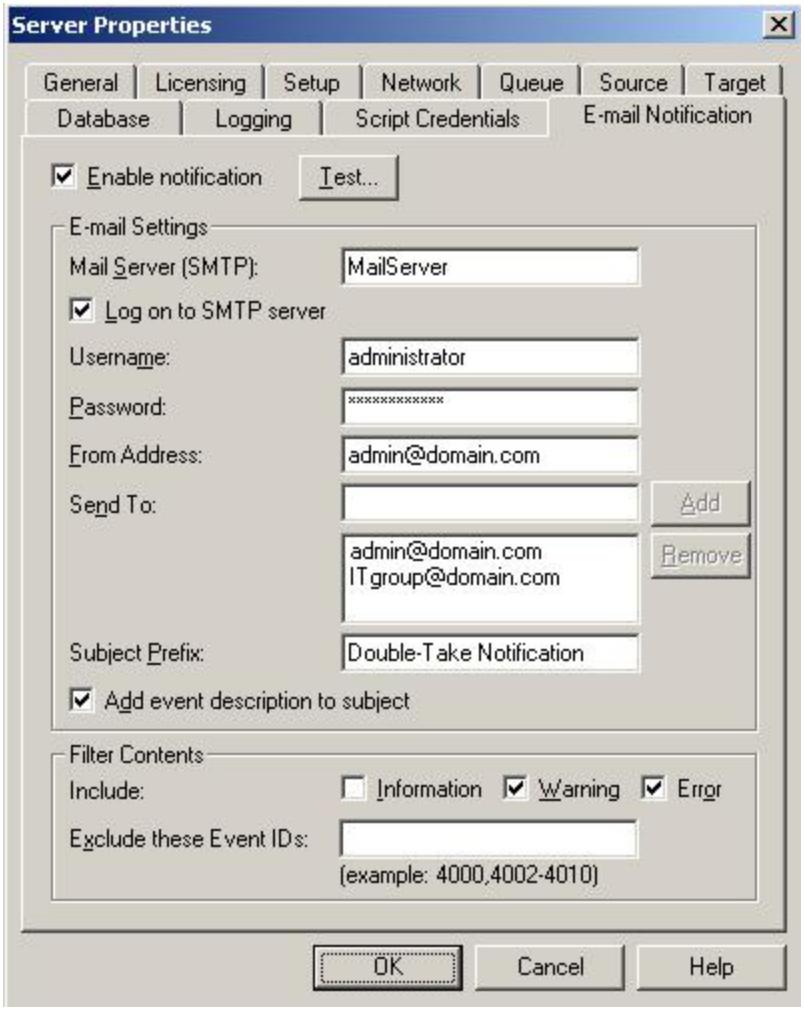

4. Select **Enable notification**.

Note: Any specified notification settings are retained when **Enable notification** is disabled.

5. Specify your e-mail settings.
   - **Mail Server (SMTP)**—Specify the name of your SMTP mail server.
   - **Log on to SMTP Server**—If your SMTP server requires authentication, enable **Log on to SMTP Server** and specify the **Username** and **Password** to be used for authentication. Your SMTP server must support the LOGIN authentication method to use this feature. If your server supports a different authentication method or does not support authentication, you may need to add the Double-Take Availability server as an authorized host for relaying e-mail messages. This option is not necessary if you are sending exclusively to e-mail addresses that the SMTP server is responsible for.
   - **From Address**—Specify the e-mail address that you want to appear in the From field of each Double-Take Availability e-mail message. The address is limited to 256 characters.
   - **Send To**—Specify the e-mail address that each Double-Take Availability e-mail message should be sent to and click **Add**. The e-mail address will be inserted into the list of addresses. Each address is limited to 256 characters. You can add up to 256 e-mail addresses. If you want to remove an address from the list, highlight the address and click **Remove**. You can also select multiple addresses to remove by Ctrl-clicking.
   - **Subject Prefix** and **Add event description to subject**—The subject of each e-mail notification will be in the format Subject Prefix : Server Name : Message Severity : Message ID : Message Description. The first and last components (Subject Prefix and Message Description) are optional. The subject line is limited to 150 characters.

     If desired, enter unique text for the **Subject Prefix** which will be inserted at the front of the subject line for each Double-Take Availability e-mail message. This will help distinguish Double-Take Availability messages from other messages. This field is optional.

     If desired, enable **Add event description** to subject to have the description of the message appended to the end of the subject line. This field is optional.
   - **Filter Contents**—Specify which messages that you want to be sent via e-mail. Specify **Information**, **Warning**, and/or **Error**. You can also specify which messages to exclude based on the message ID. Enter the message IDs as a comma or semicolon separated list. You can indicate ranges within

the list.

Note: You can test e-mail notification by specifying the options on the **E-mail Notification** tab and clicking **Test**. (By default, the test will be run from the machine where the Replication Console is running.) If desired, you can send the test message to a different e-mail address by selecting **Send To** and entering a comma or semicolon separated list of addresses. Modify the message text up to 1024 characters, if necessary. Click **Send** to test the e-mail notification. The results will be displayed in a message box. Click **OK** to close the message and click **Close** to return to the **E-mail Notification** tab.

E-mail notification will not function properly if the Event logs are full.

If an error occurs while sending an e-mail, a message will be generated. This message will not trigger an e-mail. Subsequent e-mail errors will not generate additional messages. When an e-mail is sent successfully, a message will then be generated. If another e-mail fails, one message will again be generated. This is a cyclical process where one message will be generated for each group of failed e-mail messages, one for each group of successful e-mail messages, one for the next group of failed messages, and so on.

If you start and then immediately stop the Double-Take service, you may not get e-mail notifications for the log entries that occur during startup.

By default, most virus scan software blocks unknown processes from sending traffic on port 25. You need to modify the blocking rule so that Double-Take Availability e-mail messages are not blocked.

---

6. Click **OK** to save the settings.

Full-server protection

Before you can protect your entire source, you need to select a target that is suitable to become the source, in the event of a source failure. Double-Take Availability will validate the target you select and identify any incompatibilities. Errors will disqualify the target as a suitable server. First you will need to find a compatible target, then you can establish full-server protection.

Finding a compatible target

Review the table below to determine if your server is a compatible target.

Operating system version

The source and the target must have the same operating system. For example, you cannot have Windows 2003 on the source and Windows 2008 on the target. The two servers do not have to have the same level of service pack or hotfix.

Domain

The domain of the source should be the same as the domain of the target.

Server role

The target cannot be a domain controller. Ideally, the target should not host any functionality (file server, application server, and so on) because the functionality will be removed when failover occurs.

If your source is a domain controller, it will start in a non-authoritative restore mode after failover. This means that if the source was communicating with other domain controllers before failover, it will require one of those domain controllers to be reachable after failover so it can request updates. If this communication is not available, the domain controller will not function after failover. If the source is the only domain controller, this is not an issue.

Architecture

The source and the target must have the same architecture. For example, you cannot failover a 32-bit server to a 64-bit server.

Processors

There are no limits on the number or speed of the processors, but the source and the target should have at least the same number of processors. If the target has fewer processors or slower speeds than the source, there will be performance impacts for the users after failover.

Memory

The target memory should be within 25% (plus or minus) of the source. If the target has much less memory than the source, there will be performance impacts for the users after failover.

Network adapters

You must map at least one NIC from the source to one NIC on the target. If the source has more NICs than the target, some of the source NICs will not be mapped to the target. Therefore, the IP addresses associated with those NICs will not be available after failover, unless you configure the advanced options. If there are more NICs on the target than the source, the additional NICs will still be available after failover.

File system format

The source and the target must have the same file system format. For example, an NTFS volume cannot be sent to a FAT volume.

HAL type and version

The Windows hardware abstraction layer (HAL) refers to a layer of software that deals directly with your computer hardware. The HAL type and version do not have to be identical, but they must be compatible between the source and the target. If the two are incompatible, Double-Take Availability will warn you. In that case, you must upgrade or downgrade the target.

Administrative shares

The Full-Server Failover Manager console must be able to access administrative shares on the source and the target.

Boot volume configuration

The target boot volume cannot be a dynamic disk configuration. The boot volume is the disk volume that contains the Windows operating system and supporting files. By default, the operating system files are in the \Windows folder, and the supporting files are in the \Windows\System32 folder. The boot volume might be the same volume as the system volume, but that configuration is not required.

System volume

The target must have the same system volume as the source. The system volume is the disk volume that contains the hardware-specific files that are needed to start Windows. The system volume might be the same volume as the boot volume, but that configuration is not required.

Logical volumes

There are no limits to the number of logical volumes, although you are bound by operating system limits. The source and the target must have the same number of logical volumes, and the source and the target must have

the same drive letters. For example, if the source has drives C: and D:, the target cannot have drives D: and E:. In this case, the target must also have drives C: and D:.

System path

The source and the target must have the same system path. The system path includes the location of the Windows files, Program Files, and Documents and Settings.

Double-Take Availability path

Double-Take Availability must be installed on the system path on the source and the target.

Double-Take Availability data state

The source should be from a time when the Double-Take Availability data state is good. If you are using snapshots, you may want to use a snapshot from the last good data state.

Capacity and free space

The target must have enough space to store the data from the source. This amount of disk space will depend on the applications and data files you are protecting. The more data you are protecting, the more disk space you will need. You must also have enough space on the target to process and apply the system state data.

Double-Take Availability performs several validation checks to determine if adequate disk space is available. First, the target must have enough free space on its system volume to hold the entire volume(s) (free and used) from the source. If this first validation check passes, then no additional checks are necessary. Otherwise, there must be at least enough free space on the target system volume to store the system path (including the location of the Windows files, Program Files, and Documents and Settings) from the source. If this second validation check passes, then no additional checks are necessary. If this second validation fails, Double-Take Availability will check to see if a previous failover may have been attempted. Since Double-Take Availability can reuse the disk space from a previous failover attempt, it will add the size of that data to the amount of free space available. If that is enough space for the failover, the failover will continue. If not, you will either have to select a different target or delete files on the target to free disk space.

Establishing full-server protection

Use the following instructions to establish protection for your entire source.

1. Login as a user that is a member of both the **Double-Take Admin** and local Administrators security groups.
2. Open the Full-Server Failover Manager.
3. Enter your source and target servers. You can click **Browse** when selecting either server to locate it by drilling down through your network. After you have specified a server name, enter login credentials when prompted. Once the server is selected and logged in, the **Properties** tab at the bottom of the Full-Server Failover Manager updates to display the server's properties.

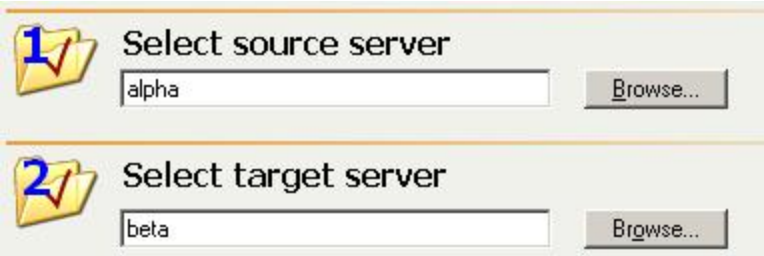

4. Optional protection settings are available but not required. If desired, configure the optional protection settings by clicking **Configure protection**.
5. You must validate that your target is compatible with your source and can stand-in if the source fails. Click **Validate** configuration. The **Validation** tab at the bottom of the Full-Server Failover Manager updates to display the validation check. Errors are designated by a white X inside a red circle. Warnings are designated by a black exclamation point (!) inside a yellow triangle. A successful validation is designated by a white checkmark inside a green circle.

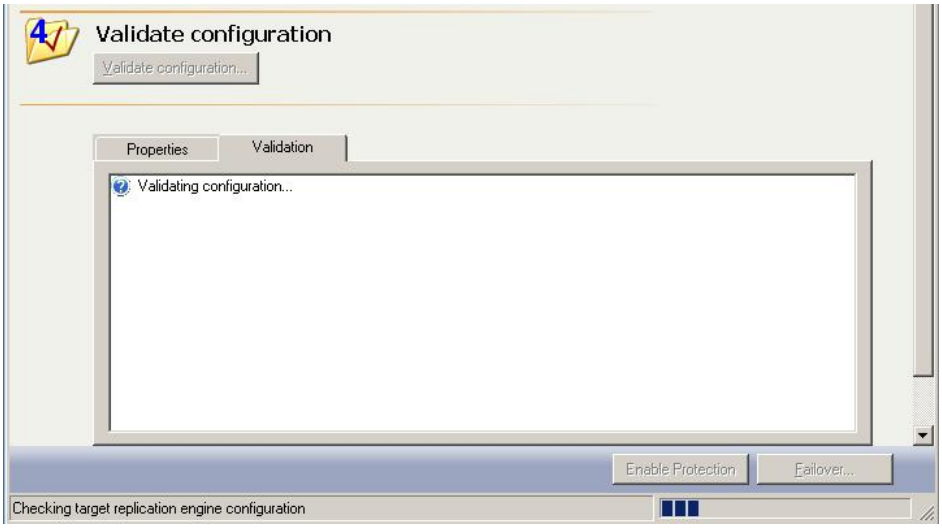

6. Double-click on any of the validation items to see details. You must correct any errors before you can enable protection. Depending on the error, you may be able to click **Fix** or **Fix All** and let Double-Take Availability correct the problem for you. For those errors that Double-Take Availability cannot correct automatically, you will need to modify the target to correct the error, or you can select a different target. You must revalidate the selected servers until the validation check passes without errors.
7. Once the validation check passes without errors, click **Enable Protection** to begin monitoring.

Note: After protection is enabled, a **View Configuration** link will display the optional protection settings in read-only mode. If you need to modify any of the optional protection settings, you will have to disable your protection, modify the optional protection settings, and then re-enable protection.

If you have a full-server protection connection established, do not create any other Double-Take Availability connections from or to the source or target.

If you are using a cluster, you must manually alter the disk signature before or after failover. See the Microsoft Knowledge Base article 280425 for details on how to change the disk signature. You can automate the disk signature alteration as part of failover using the pre- or post-failover scripts.

Optional full-server protection settings

Optional protection settings are available when configuring a full-server protection connection, but they are not required. If you want to configure the optional settings, make sure you have a valid source and target specified, then click **Configure protection** from the main Full-Server Failover Manager page.

You have the following optional configuration settings available.

- Including and excluding data to be protected
- Stopping services on the target when protection is enabled
- Taking snapshots of the target
- Configuring failover monitoring and processing
- Mapping network configuration on the target for post-failover
- Routing data transmissions
- Mirroring data
- Compressing data

Including and excluding data to be protected

1. Make sure you have a valid source and target specified, click **Configure protection** from the main Full-Server Failover Manager page, and then select the **Protection** tab.
2. Specify that data you want to include and/or exclude.
    - **Volumes to include**—Select the volumes that you want to protect. You must have the same volumes on the source and target. You cannot deselect the boot volume. If you deselect other volumes, you will not be protecting the entire source server.
    - **Directories to exclude**—Select the directories that you want to exclude. If you exclude any directories, you will not be protecting the entire source server.
        - **Add**—Click **Add** to specify a directory to exclude. Enter the name of the directory or click **Browse** to search for the directory path. Specify **Recursive** if sub-directories of the listed directory should be excluded also. Specify **Non-recursive** if sub-directories of the listed directory should be included. Click **OK** to add the directory to the list of Directories to exclude. Repeat this process to add multiple directories.
        - **Edit**—Highlight a directory in the list and click **Edit** to modify the directory definition. After modifying the directory, click **OK** to save the changes.
        - **Remove**—Highlight a directory in the list and click **Remove** to delete the directory definition. The directory will no longer be excluded.
3. Click **OK** to save the settings.

Stopping services on the target when protection is enabled

1. Make sure you have a valid source and target specified, click **Configure protection** from the main Full-Server Failover Manager page, and then select the **Protection** tab.
2. Double-Take Availability determines what services are currently running on the target. You can specify which services you want to keep running and those services you want to stop when you enable your protection. Move the services between the **Services to stop** and **Services to leave running** lists by using the double arrows. Select **Show critical services** to see the list of critical services that will remain running on the target. The critical services are displayed in a lighter colored, italics font. The critical services cannot be moved from the running list.
3. Click **OK** to save the settings.

Taking snapshots of the target

1. Make sure you have a valid source and target specified, click **Configure protection** from the main Full-Server Failover Manager page, and then select the **Protection** tab.
2. A snapshot is an image of data taken at a single point in time. Snapshots allow you to view files and folders as they existed at points of time in the past, so you can, for example, recover from cases where corrupted source data was replicated to the target. By default, Double-Take Availability takes periodic snapshots of the data on the target. When failover is triggered, you can use the live target data at the time of failover or you can failover to a snapshot of the target data. Specify how you want to handle snapshots.
    - **Enable periodic snapshots**—Enable snapshots if you want to be able to use them at failover time. If snapshots are disabled, the live data on the target at failover time will be used.
    - **Snapshot Interval**—By default, snapshots of the target data are taken every 60 minutes. If desired, increase or decrease the interval between snapshots.
    - **Start now**—If you want to start taking snapshots immediately after the full-server connection is established, select **Start now**.
    - **Start at**—If you want to start taking snapshots at a specific data and time, select **Start at** and specify the date and time parameters.
3. Click **OK** to save the settings.

Configuring failover monitoring and processing

1. Make sure you have a valid source and target specified, click **Configure protection** from the main Full-Server Failover Manager page, and then select the **Failover** tab.
2. Specify how you want to handle failover monitoring and processing.
    - **Manual intervention required**—By default, you will be notified when a failover is necessary, but the failover process will not start until you manually initiate it. If you disable intervention, failover will automatically start when a failure is detected.
    - **Monitor Interval**—By default, the target checks to see if the source is online every five (5) seconds. The source responds back to the target when it receives one of these checks. If desired, increase or decrease the interval between checks.
    - **Missed intervals**—By default, the target can miss five (5) responses from the source before assuming the source has failed. If desired, increase or decrease the number of responses that can be missed before the source is identified as failed.

    Note: Together, the **Monitor interval** and **Missed intervals** settings determine the total time before failover would be triggered. For example, five missed checks every five seconds would be 25 seconds to trigger failover. To achieve shorter delays before failover, use lower values. To achieve longer delays before failover, choose higher values.

    - **Monitor type**—Double-Take Availability monitors the source by testing for a Double-Take Availability network response via ICMP pings or a Double-Take service response. Select the type of monitoring you want to perform.
    - **Pre-Failover Script** and **Post-Failover Script**—If you want to execute a script on the target before failover (pre-failover) begins or after failover has occurred (post-failover), specify the path to the script on the target. You can also search for the script by clicking **Browse**. The script will be processed using the same user account that is configured to run the Double-Take service. Scripts may contain any valid Windows command, executable, batch, or script file. The **Pre-Failover Script** will be executed as soon as the failover process is initiated. The **Post-Failover Script** will be executed after the failover process ends and the target has been rebooted.
3. Click **OK** to save the settings.

Mapping network configuration on the target for post-failover

1. Make sure you have a valid source and target specified, click **Configure protection** from the main Full-Server Failover Manager page, and then select the **Failover** tab.
2. Specify how you want to handle network configurations on the target after failover.
    - **Apply source network configuration to the target**—If you select this option, the source IP addresses will be failed over to the target. The IP addresses associated with each NIC on the source will be mapped to the NIC you specify on the target. If the source has more NICs than the target, some of the source NICs will not be mapped to the target, or the target NICs may need to be used for more than one source NIC. If there are more NICs on the target than the source, the additional NICs will still be available after failover. If you are using a LAN environment, you should select this option.
    - **Retain target network configuration**—If you select this option, the source IP addresses will not be failed over to the target. The target will retain all of its original IP addresses. If your target is on a different subnet (typical of a WAN environment), you should select this option. The following options are available if you want to update DNS.
    - **DNS Server**—The source's primary DNS server is selected by default. If desired, you can select a different DNS server.
    - **Source Server IP**—Select the IP address from the source that you want to remap to an IP address on the target.
    - **Target Server IP**—Select an IP address on the target that will be replace the source IP address in DNS.
    - **Username**—Specify a user that has privileges to access and modify DNS records. The account must be a DNS Admin for the domain in which the DNS server resides. You can enter a user name for a different domain by entering a fully qualified user name. The fully qualified user name must be in the format domain\username or username@domain. If you enter a non-qualified name, the DNS domain will be used by default. The domain name is obtained from the DNS server name, provided that reverse lookup in DNS is enabled.
    - **Password**—Enter the password that is associated with the specified user name.
3. After you have entered the DNS information, click **Test** to validate that DNS is configured correctly and that the specified credentials are sufficient to update DNS. When the DNS configuration is complete, click **OK** to save your entries and return to the **Configure Protection** window.

Note: If you are protecting a cluster node and select **Retain target network configuration**, the virtual IP addresses associated with the cluster will still be failed over. The only IP address that will not be failed over is the physical IP address of the cluster node.

4. Click **OK** to save the settings.

Routing data transmissions

1. Make sure you have a valid source and target specified, click **Configure protection** from the main Full-Server Failover Manager page, and then select the **Advanced** tab.
2. By default, Double-Take Availability will select the default route for transmissions. If desired, select a different IP address on the target that will be used for full-server connection transmissions.
3. Click **OK** to save the settings.

Mirroring data

1. Make sure you have a valid source and target specified, click **Configure protection** from the main Full-Server Failover Manager page, and then select the **Advanced** tab.
2. Specify how you want to mirror data.
    - **Mirroring**—Select the type of Double-Take Availability mirroring process you want to perform. A **Full** mirror will transmit all files from the source to the target. A **Checksum** mirror will transmit only the blocks of data that are different between the source and target.
    - **Estimate Replication Set size based on volume**—The size of the replication set is used to determine mirroring remaining calculations. If this option is enabled, the replication set size is based on the volume size, which is a faster calculation. If disabled, the replication set size is based on the selected data, which can be a slower calculation. The mirror will remain in an initializing state until the calculation is complete.
3. Click **OK** to save the settings.

Compressing data

1. Make sure you have a valid source and target specified, click **Configure protection** from the main Full-Server Failover Manager page, and then select the **Advanced** tab.
2. Compression allows you to reduce the amount of bandwidth needed to transmit data from the source to the target. The data is compressed before being transmitted and then is uncompressed before it is written on the target. Typically, compression is used in WAN environments, but not in LAN environments. If desired, enable compression and select the level of compression that you want to use. All connections to the same target will have the same compression settings.
3. Click **OK** to save the settings.

Using NAT or firewalls with full-server workloads

If your source and target are on opposite sides of a NAT or firewall, you will need to configure your hardware to accommodate full-server workload communications. You must have the hardware already in place and know how to configure the hardware ports. If you do not, see the reference manual for your hardware.

- Full-server ports
- Microsoft Windows ports
- Hardware ports

Full-server ports

By default, Double-Take Availability uses port 6320 for all communications. To verify or modify the ports, you must use the Replication Console.

1. Open the Replication Console.
2. Locate your servers in the server tree in the left pane of the Replication Console.
3. Right-click the server in the left pane of the Replication Console and select **Properties**.
4. On the **Network** tab, verify or modify the **Communications Port** as needed. All servers and clients in your full-server workload protection must be using the same port.

Double-Take Availability uses ICMP pings to monitor the source for failover. A failover monitor will not be created if ICMP is blocked (although the data and system state will still be protected). You should configure your hardware to allow ICMP pings between the source and target. If you cannot, you will have to monitor for a failure using the Double-Take Availability replication service. Use the **Monitor Type** option on the configuration **Failover** tab to set your failover monitoring method.

Microsoft Windows ports

Full-server workload protection uses WMI (Windows Management Instrumentation) which uses RPC (Remote Procedure Call). By default, RPC will use ports at random above 1024, and these ports must be open on your firewall. RPC ports can be configured to a specific range by specific registry changes and a reboot. See the Microsoft Knowledge Base article 154596 for instructions.

Full-server workload protections also rely on other Microsoft Windows ports.

- Microsoft File Share uses ports 135 through 139 for TCP and UDP communications.
- Microsoft Directory uses port 445 for TCP and UDP communications.

These ports must be open on your firewall. Check your Microsoft documentation if you need to modify these ports.

Hardware ports

You need to configure your hardware so that the full-server ports and Microsoft Windows ports are open. Since communication occurs bidirectionally, make sure you configure both incoming and outgoing traffic.

There are many types of hardware on the market, and each can be configured differently. See your hardware reference manual for instructions on setting up your particular router.

Application protection

Protecting an application consists of three main tasks - configuring the protection, validating the protection, and enabling the protection. The processes are similar for the different application types, however, there are some variances. Optional protection settings are available when configuring application protection, but they are not required.

Protecting an application

1. Open the Application Manager.
2. Verify the **Setup** tab is selected and then from the **Tasks** list on the left pane, select the type of application you want to protect.

Note: The fields in the Application Manager console will vary depending on the type of application you are protecting.

---

3. Application Manager will automatically identify the root domain where the Application Manager is running and populate the **Domain Name** field. If necessary, change the domain name to a trusted root domain that the Application Manager console can connect to. If prompted, enter security credentials with administrator privileges for the domain.

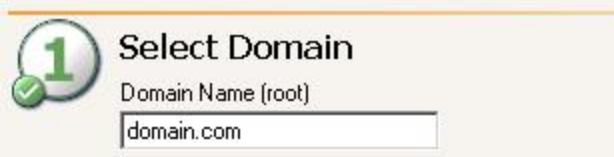

Note: Domain names must include a suffix, such as .com, .corp, .net, and so on.

Exchange Note: If you are protecting Exchange, the domain must be the root of the forest domain because that is where all Exchange server objects reside, even if the Exchange server is a member of a child domain.

---

4. Application Manager will automatically attempt to populate the **Source Server** and **Target Server** lists with any servers in the specified domain that are running the application you are protecting. Select your source and target servers.

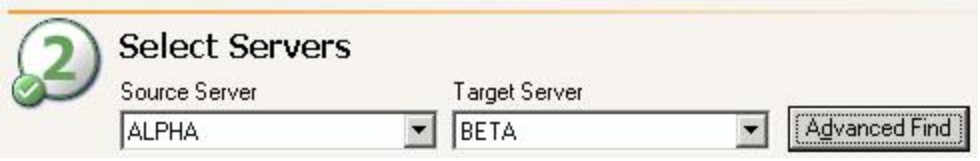

Note: If you have previously used this source and target pair, you will be prompted to reuse the previous configuration. If you select **Yes**, your previous configuration settings will be used. If you do not want to use the previous configuration settings (perhaps the source or target configuration has changed since you configured the connection), select **No** to use the default configuration settings.

If you select a source that is currently unavailable, you will be prompted to select the target first. When you select the target then the source, you may get a failover prompt or the same source is unavailable prompt. This will depend on if a failover condition has been met according to the original failover configuration.

You cannot protect a server if it is already functioning as a target.

If no servers are populated in the lists (perhaps the server you need is in a child domain), click **Advanced Find** to add servers to the lists. **Advanced Find** is not available for all application protections. See Managing servers for more details on **Advanced Find**.

Server names must be 15 characters or less.

You will be unable to configure protection when your environment between the Application Manager and the source or target contains a NAT or certain VPN configurations. This is due to WMI limitations. Contact technical support for instructions to configure protection manually.

Exchange Note: The target you select must be in the same Exchange administrative group as the source.

If you are protecting Exchange 2003 and it is running in mixed mode, the first installed Exchange virtual server contains the MTA (Message Transfer Agent) resource that is needed to communicate with versions prior to Exchange 2003. If you do not failover all Exchange virtual servers, then any user who is in a different mail store than the first one may not be able to route mail.

If you are protecting multiple Exchange virtual servers, you can configure multiple like-named cluster protection connections, or you can failover multiple Exchange virtual servers to pre-existing Exchange virtual servers on the target. In this case, select the target Exchange virtual server.

If you are protecting Exchange in a like-named cluster scenario, select the same server for the source and target. The target server name will automatically be appended with the suffix like-named. Enter the requested information in the **Like-named Cluster Setup** dialog box.

- **Target Cluster**—Enter the name of one of the target nodes, then click **Connect**.
- **Network**—Select the target NIC that can accommodate a new IP address.
- **IP Address**—Enter a new IP address for the target to use when it stands in for the source.
- **Subnet Mask**—Enter the subnet mask to use for the new IP address.
- **Storage Resources**—The Application Manager will automatically select the required storage resources on the target, provided that they exist. Verify that the drive letters where Exchange data is located are selected. You cannot deselect a storage resource that exists on both the source and target. If the drive letters on the source and target do not match, then not all required data will be selected automatically. You will need to select it manually. The selected storage resources must be in the same group.

SharePoint Note: If you are protecting SharePoint, enter the **SharePoint Front-End Web Server** and click **Get Config** to populate the **Source Server** list.

---

5. If prompted when selecting a server, provide login credentials.

Note: You can enter a user for a different domain by entering a fully qualified name in the format domain\username or username@domain. If you enter a

non-qualified name, the DNS domain from the DNS server will be used.

The login account must be a member of the local administrators security group.

The login account must be a member of the Double-Take Admin security group.

If you will be configuring DNS failover, the login account must be a member of the domain DnsAdmins security group.

Exchange Note: The login account must be an Exchange Full Administrator at the organizational level, as delegated via the Exchange System Manager at the user level or have delegated rights via the Application Manager Delegate Rights feature (select **Tools**, **Delegate Rights**). Rights must be delegated to a specific user and not the group the user belongs to in order for the Application Manager to recognize them.

The login account must have rights to manage Exchange in order to query and modify the Exchange Active Directory objects.

If Exchange is on a cluster, the login account must be a member of the Cluster Administrators security group on each node. Additionally, the same cluster service account should be used for both the source and target.

SQL Note: The login account must be assigned the System Administrator role on the SQL server in order to query and administer SQL.

SharePoint Note: The login account must be assigned the System Administrator role on the SQL server in order to query and administer SQL.

The login account must be the same farm administrator account used when installing and configuring SharePoint.

The login account must have local administrator rights on the source and target front-end web servers and the source and target back-end SQL servers.

BlackBerry Note: If the login account is the BesAdmin account, you will need to add that account to the Double-Take Admin group on the source and target servers.

6. Optional protection settings are available but not required. If desired, configure the optional protection settings by clicking **Configure**.
7. Once you have finalized your protection settings, you need to validate your configuration by clicking **Validate**. Monitor the status of the validation process in the status bar at bottom of the Application Manager.
8. When the validation is complete, the status progress indicator is removed. Errors are designated by a white X inside a red circle. Warnings are designated by a black exclamation point (!) inside a yellow triangle. A successful validation is designated by a white checkmark inside a green circle. Unknown items are a white question mark inside a blue circle. Double-click on any of the validation items to see details. Application Manager can automatically fix errors and warnings marked with a yellow gear. Highlight a single item and click **Fix** to have Application Manager fix that one issue. Click **Fix All** to have Application Manager correct all of the issues. You must correct any errors, and ideally any warnings, before you can enable protection.

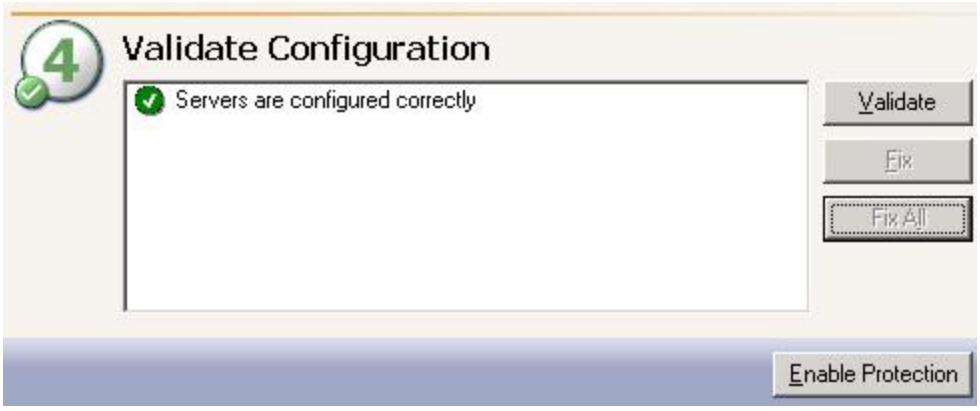

Note: If you are using DNS failover and did not enter DNS credentials under the optional protection settings, you will be prompted for credentials that can

access and modify DNS records.

If you validate a source and target pair that is already in a **Protected** state and the validation detects issues with the target, **Fix** and **Fix All** will be disabled. You must disable the protection, fix the issue, then re-enable protection.

SQL Note: If you are using Database Only mode and the database is online on the target, the database cannot be taken offline on the target by using **Fix** if it has a SQL Server replication publication. The publication will have to be deleted using the SQL Server management tools before the database can be taken offline.

For SQL 2000 servers, Application Manager may hang when rerunning validation after disabling protection in the same session. To work around this issue, disable the protection, stop and restart the Application Manager, then validate or enable protection.

---

9. If you are protecting BlackBerry, a Complete BlackBerry Protection dialog will appear after the validation is complete. Select the name of the Exchange server under step 2 of the dialog box and you will be able to complete the application protection process again in order to protect your Exchange server.
10. Once the validation passes without errors, click **Enable Protection**. View the status of the protection on the Monitor tab.

Note: If you modify your source server configuration on the source server, for example, adding a new storage group or database, you must disable protection, run validation and fix any issues, then re-enable protection to apply the changes.

If your application protection is in a cluster environment, you should not move any resources from one cluster group to another once protection is established.

If you close Application Manager prior to enabling protection, your configuration changes will not be saved. You must enable protection in order to save your configuration settings.

After you have enabled protection, do not use the Failover Control Center client to edit your failover configuration. Doing so will force a failover. If this occurs, cancel the failover prompt and then disable and re-enable monitoring using Application Manager.

Exchange Note: If you need to protect data that is stored on a non-mailbox server role, for example SMTP queue data, you will need to configure protection for that data separately. In addition, you may need to manually update the DNS setting for the client access server to point to the target site.

If your source has a public store that has a non-MAPI Owning Tree, the folders within this tree may not be updated with current information and therefore unavailable to the users.

If you are protecting Exchange 2007 and your environment has more than one domain controller and the domain controller that Exchange is using will change during the failover process, stores may not mount during failover. To work around this issue, either modify the failover script to use the domain controller that Exchange uses when the Information Store service is started or mount the stores manually and run the Replace replicas PowerShell call from the failover script to update the PF replicas (if applicable).

If you are protecting Exchange on a cluster and the target cluster has more than one IP address resource for the virtual Exchange server you may experience the following issue. If the one of the IP addresses is not routable from the source server and that IP address was created before any of the routable IP address resources, the Application Manager will fail to enable protection. To enable protection, you will need to delete the non-routable IP address resource(s), re-create them, and then re-add them as dependencies on the network name resource for the virtual server.

If you are protecting Exchange in a like-named cluster scenario, Application Manager will create four resources on the target cluster: two generic script resources, an IP address resource, and a temporary name resource. The temporary name resource will be the Exchange virtual server with the suffix _LN. Application Manager uses the temporary name resource for the connection between the source and target clusters.

SQL Note: If you are protecting SQL on a cluster and the target cluster has more than one IP address resource for the virtual SQL Server server you may experience the following issue. If the one of the IP addresses is not routable from the source server and that IP address was created before any of the routable IP address resources, the Application Manager will fail to enable protection. To enable protection, you will need to delete the non-routable IP address resource(s), re-create them, and then re-add them as dependencies on the network name resource for the virtual server.

Optional application protection settings

Optional protection settings are available when configuring application protection, but they are not required. If you want to configure the optional settings, make sure you have a valid domain and servers specified, then click **Configure** from the main Application Manager page.

You have the following optional configuration settings available.

- Configuring failover processing
  - Configuring DNS failover
  - Configuring identity failover
- Configuring failover monitoring
- Taking snapshots of the target
- Application connection settings
  - Routing data transmissions
  - Protection configuration
    - Configuring Exchange storage group protection
    - Configuring SQL database protection
    - Configuring file share protection
    - Configuring BlackBerry database protection
    - Configuring SharePoint database protection
  - Mirroring data
- Application advanced settings
  - Configuring the replication set
  - Configuring scripts
  - Configuring Active Directory
  - Configuring items to failover
  - Configuring default connection parameters

Configuring failover processing

1. Make sure you have a valid domain and servers specified, click **Configure** from the main Application Manager page, and then select the **Failover** tab.

Note: The fields on the **Failover** tab will vary depending on the type of application you are protecting.

---

2. Specify the **Failover Type** which is the name resolution method that will be used to redirect users to the target in the event the source fails. Ideally, you should use DNS failover to reduce downtime and avoid server name or IP address conflicts. However, you may want to select identity failover if access to the domain controller or DNS server is not available, the time required to propogate DNS updates is unacceptable, or your end-users are configured to connect to an IP address and not a server name. For cluster environments, DNS failover is the only option available.

   - **DNS Failover**—With DNS failover, DNS records associated with the source will be updated to point to the target's IP address. This includes A, MX, and PTR-type DNS records. Clients will resolve the source server name to the target server's name and IP address at failover. Click **Configure** to specify the DNS failover settings.
   - **Identity Failover**—With identity failover, the target will assume the source server's name and IP address. This option can potentially cause name and/or IP address conflicts if the source is brought back online while the target is standing in. Click **Configure** to specify the identity failover settings.
3. Application Manager automatically determines the appropriate application services or resources to start and stop based on the application you are protecting, your operating system, and your application configuration. If necessary, modify the list of services or resources.

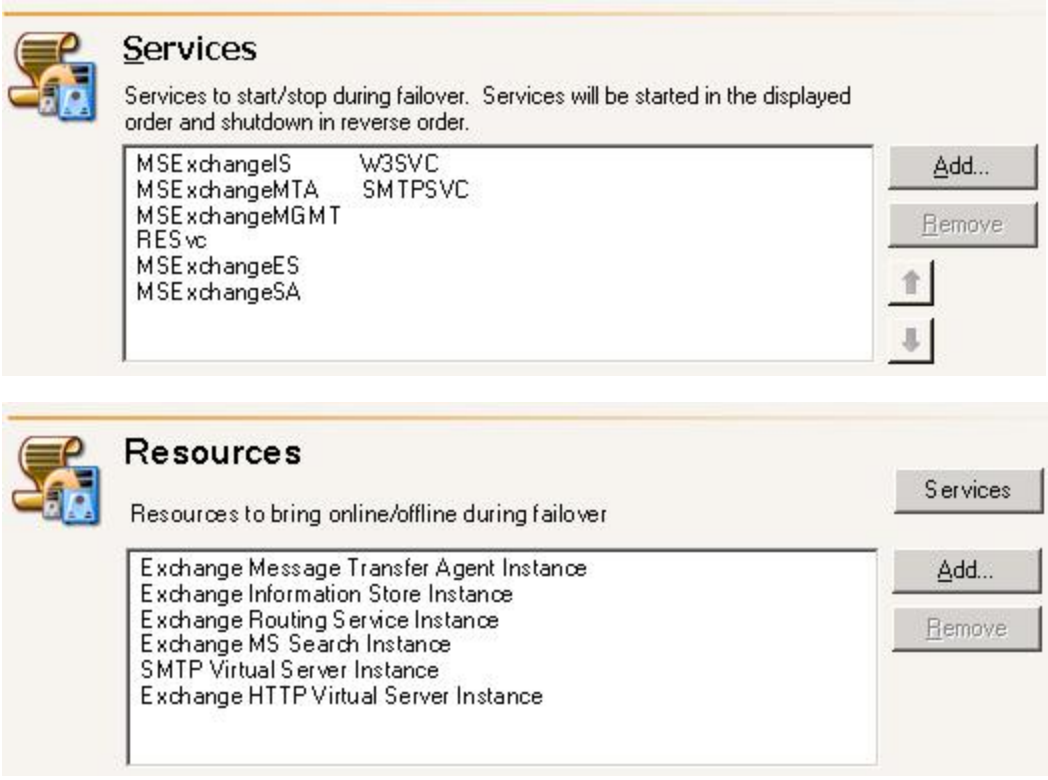

a. Click **Add** to insert a service or resource into the list. Specify the service or resource name and make sure that **Service must be stopped on target** is enabled so that replication can update files on the target.
b. Click **Remove** to remove a service or resource from the list. The services you can remove depend on the application you are protecting. For example, if you are protecting Exchange, you can only remove services or resources that you have manually added.
c. If you are configuring your services, highlight a service and click the up or down arrow to reorder the list. The services will be stopped and started in the order displayed.
d. If you are using a cluster source and standalone target configuration, you can toggle between the services and resource configuration by clicking the **Services** button.
e. If you are protecting SharePoint, you can add or remove the SharePoint services to the list of SQL services by using the appropriate **Add SharePoint Services** and **Remove SharePoint Services** buttons.

4. Click **OK** to save the settings.

Configuring DNS failover

Application Manager will automatically determine default DNS failover settings. Use the following instructions to modify the DNS failover settings.

1. The **DNS Server** list contains all DNS IP addresses for the source and target servers. The label after the IP address indicates if the DNS IP address belongs to the source, target, or both. To additional DNS servers to the list, enter an IP address into the **DNS Server** field and click **Add**. To remove an IP address from the list, highlight the address and click **Delete**.

Note: If you want to set the primary DNS server that Double-Take Availability will use during failover, you can specify **Client DNS Server**. This option is only available if you have launch the Application Manager using the command line advanced option.

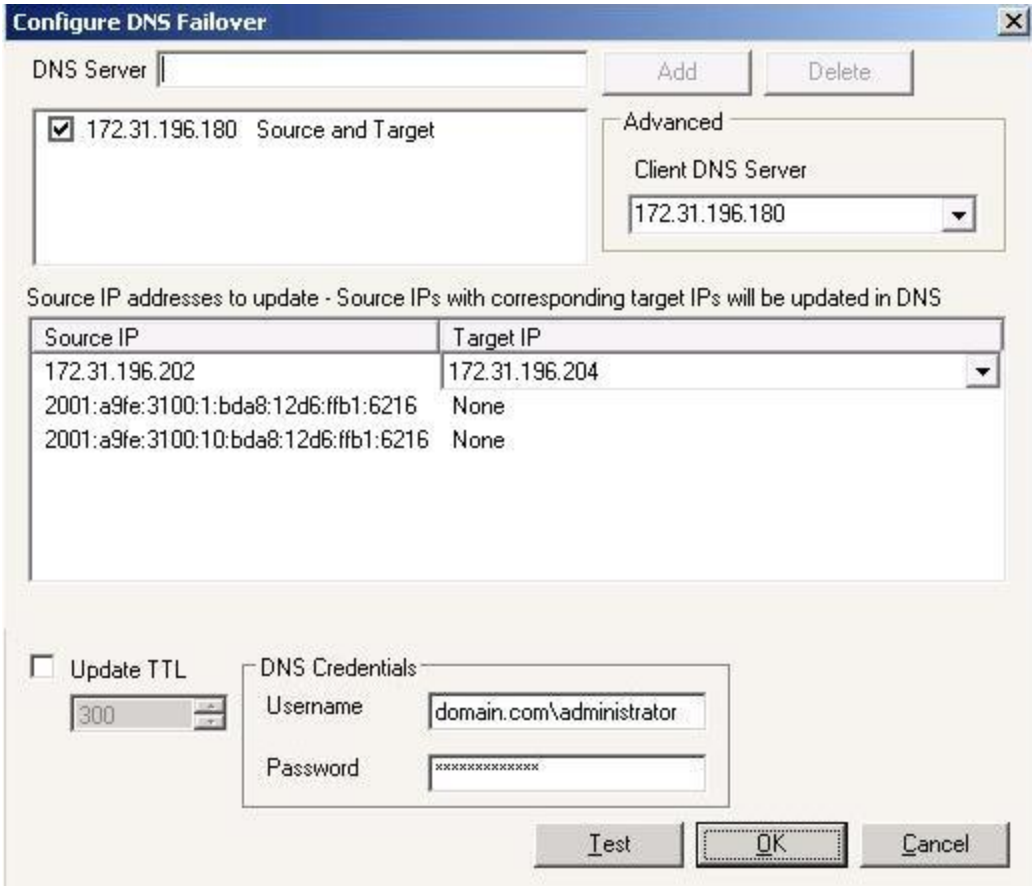

2. Under **Source IP addresses to update**, map a source IP address to a target IP

address for DNS updates.

Exchange Note: If you are protecting Exchange and one or more IP addresses are configured for the SMTP virtual server on the target, the first IP address will be the default target IP address for all source IP addresses.

---

3. You can specify the length of time, in seconds, that the source's DNS A records are cached in the Time to Live. Enable **Update TTL** and specify a number of seconds. Ideally, you should specify 300 seconds (5 minutes) or less.

Note: In order to update the Time to Live, the machine where the Application Manager is running must be able to connect to the DNS server through WMI. If it cannot, the Time to Live record will not be updated and the Application Manager will return an error that the RPC server is unavailable.

---

4. Specify DNS Credentials and specify a user that has privileges to access and modify DNS records. The user must be a member of the DNSAdmins gruop for the domain where the DNS server resides. You can enter a user for a different domain by entering a fully qualified name in the format domain\username or username@domain. If you enter a non-qualified name, the DNS domain from the DNS server will be used.
5. Once your DNS failover settings are configured, click **Test** to validate your configuration. If you have any issues with your configuration, review the following DNS information.
    - The DNS zone **Dynamic updates** should be set to **Secure only**. Otherwise, you must disable dynamic registration on the source server in order to prevent the source from reclaiming its DNS record.
    - DNS reverse lookup should be enabled For more information on enabling reverse lookup, see your Microsoft documentation.
    - If you are running Windows Server 2000 on the primary DNS server and are hosting zones or domains that contain source and/or target records, you must have the DNS WMI Provider installed on that DNS server.
    - If a hosts file entry for the source server exists on end-user machines, errors may occur during a failover and failback.

- If you have a NIC designated only for Double-Take Availability traffic, DNS registration for that NIC should be disabled.
- If you are protecting Exchange and your server is using a public IP address to receive e-mail, you will have to change the public advertised DNS MX record to reflect the target IP address. Consult your service provider for instructions.
- If you are protecting Exchange and you want to allow external e-mail to be delivered to the target server when the source is unavailable, you should create an additional external MX record for the target server. The target MX record should have a lower priority than the source. Refer to your router or firewall documentation for more information.
- In a cluster environment with more than one NIC on a server/node, if the heartbeat network IP address is at the top of the binding order, a socket error will occur when you attempt to select or configure the server. To avoid this issue, change the binding order on all source and target servers/nodes so the domain IP address is at the top of the order. Each server must be rebooted for the change to take effect.

6. Click **OK** to save your settings.

Configuring identity failover

Application Manager will automatically determine default identity failover settings. Use the following instructions to modify the identity failover settings.

1. Under **Source IP address to failover**, map a **Source IP** address to a **Target NIC**. The target NIC will assume the source IP address during failover. The **Target IP addresses** list displays the IP address(es) of the selected **Target NIC**.

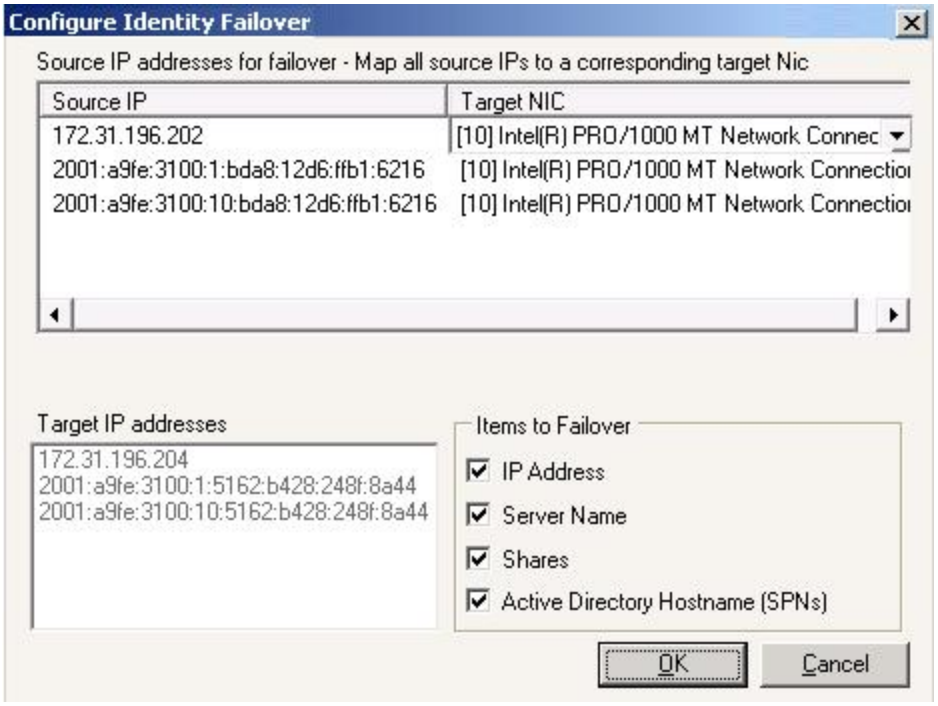

2. The default selected items under **Items to Failover** will depend on the application you are protecting. Enable or disable if you want to failover the source's IP addresses, server name, file shares, and/or Active Directory host name.

Note: If your source and target are on different subnets, you should not failover the IP address.

Exchange Note: If you are protecting Exchange, do not failover the Active Directory host name.

SQL Note: If you are protecting SQL, do not failover the server name.

File Server Note: You cannot failover file shares from a parent to child domain.

3. Click **OK** to save your settings.

Configuring failover monitoring

1. Make sure you have a valid domain and servers specified, click **Configure** from the main Application Manager page, and then select the **Monitoring** tab.
2. To enable the target to monitor the source for a failure, enable **Active Monitoring Enabled**. If this option is disabled, you will need to manually monitor the source and initiate failover if the source fails.
3. To help you better control when failover occurs, enable **Manual Intervention Required**. If this option is disabled, failover will occur immediately when a failover condition is met.
4. Specify the following **Monitor Settings**.

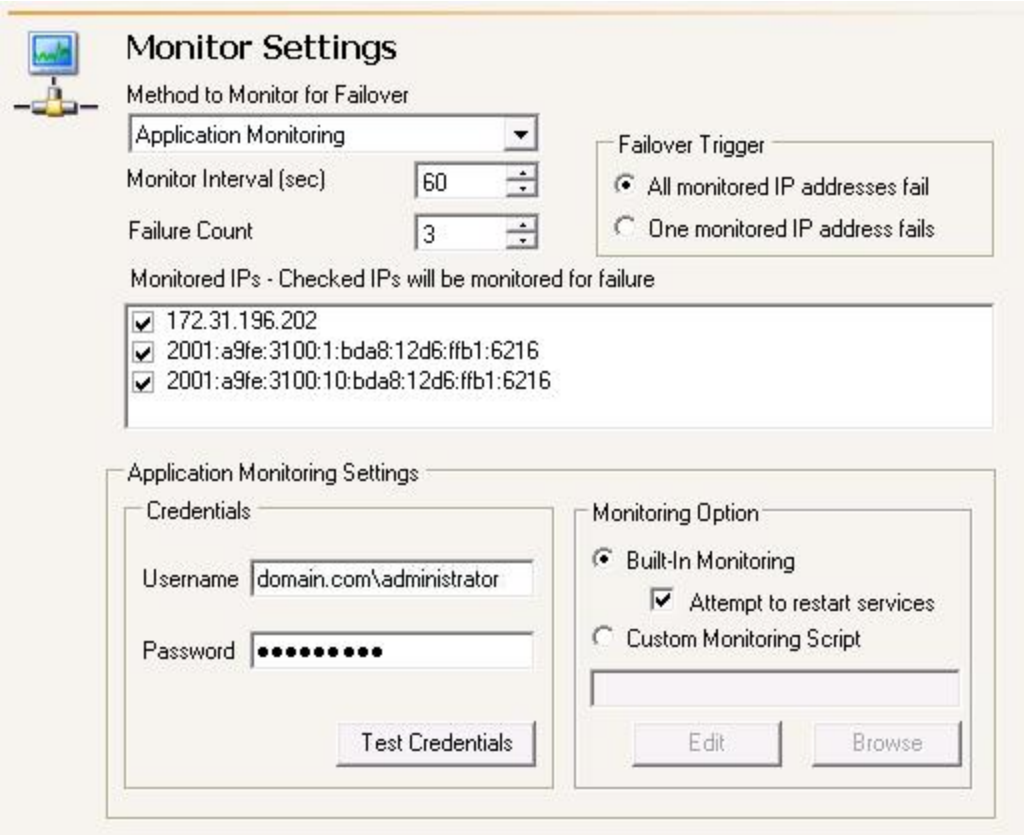

- **Method to Monitor for Failover**—You have three choices for having the target monitor the source for a failure.
- **Network Access**—ICMP pings are used to determine if the source is online. Network devices, such as firewalls or routers, must have ICMP pings unblocked in order to use this option.
- **Replication Service**—The Double-Take service on the target sends a UDP request to the source, which replies immediately to confirm it is online. Use this method if ICMP pings are blocked.

- **Application Monitoring**—In non-cluster environments, you can monitor an application via a monitoring script. The script can be WMI, PowerShell, Visual Basic, or JScript.
- **Monitor Interval**—This setting identifies the number of seconds between the monitor requests sent from the target to the source to determine if the source is online. The default setting depends on the **Method to Monitor for Failover**.
- **Failure Count**—This setting is the number of monitor replies sent from the source to the target that can be missed before assuming the source machine has failed. The default setting depends on the **Method to Monitor for Failover**.

Note: To achieve shorter delays before failover, use lower **Monitor Interval** and **Failure Count** values. This may be necessary for IP addresses on servers that must remain available and responsive at all times. Lower values should be used where redundant interfaces and high-speed, reliable network links are available to prevent the false detection of failure. If the hardware does not support reliable communications, lower values can lead to premature failover. To achieve longer delays before failover, choose higher values. This may be necessary for IP addresses on slower networks or on a server that is not transaction critical. For example, failover would not be necessary in the case of a server restart. Additionally, in a cluster environment, you must take into account the time it will take for a virtual server to failover between nodes.

- **Monitored IPs**—Mark the IP addresses on the source that you want the target to monitor.
- **Failover Trigger**—Specify if you want failover to be triggered when all monitored IP addresses fail or if only one of the monitored IP addresses fail.
- **Application Monitoring Settings**—If you selected **Application Monitoring** as your **Method to Monitor for Failover**, you need to configure the script that will monitor your application.
- **Monitoring Option**—You have a choice of two monitoring methods.
- **Built-In Monitoring**—Double-Take Availability will monitor the application using WMI. Specify if you want Double-Take Availability to **Attempt to restart services** if the application services are not responding.
- **Customer Monitoring Script**—Click **Browse** to locate your own custom

script that will monitor the application. The custom script can be PowerShell, Visual Basic, or JScript. Once your script is identified, you can click **Edit** to open a text editor to view or modify the script.

Note: A sample Visual Basic script for application monitoring is available in the \Application Manager\Samples subdirectory where Double-Take Availability is installed.

If you are using a PowerShell script, PowerShell must be installed on the target.

- **Credentials**—Specify a user with access to run the application monitoring script on the target. The user name must be a fully qualified name in the format domain\username or username@domain. If you are using the **Built-In Monitoring** script, the user must have full WMI access to the CIMV2 namespace. By default, the administrative group on Windows 2008 has full WMI access. If you are using Windows 2008, the user must be allowed through DCOM and User Access Control.

5. Click **OK** to save the settings.

Taking snapshots of the target

1. Make sure you have a valid domain and servers specified, click **Configure** from the main Application Manager page, and then select the **Snapshot** tab.

   Note: Snapshots are not available in cluster environments.

   ---

2. A snapshot is an image of data taken at a single point in time. Snapshots allow you to view files and folders as they existed at points of time in the past, so you can, for example, recover from cases where corrupted source data was replicated to the target. By default, Double-Take Availability takes periodic snapshots of the data on the target. When failover is triggered, you can use the live target data at the time of failover or you can failover to a snapshot of the target data. Specify how you want to handle snapshots.

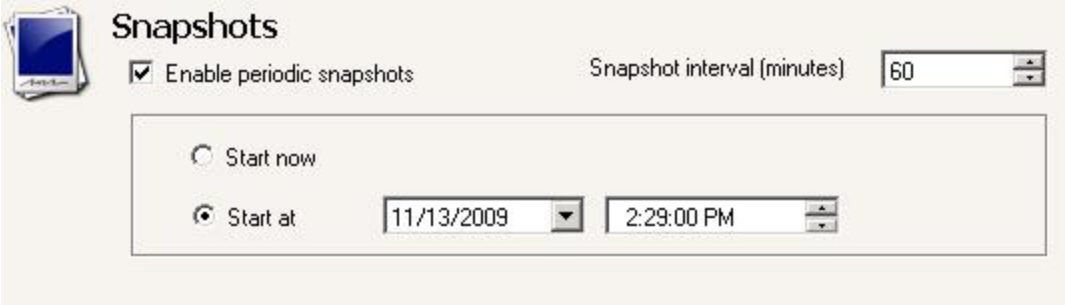

   - **Enable periodic snapshots**—Enable snapshots if you want to be able to use them at failover time. If snapshots are disabled, the live data on the target at failover time will be used.
   - **Snapshot Interval**—By default, snapshots of the target data are taken every 60 minutes. If desired, increase or decrease the interval between snapshots.
   - **Start now**—If you want to start taking snapshots immediately after the application connection is established, select **Start now**.
   - **Start at**—If you want to start taking snapshots at a specific data and time, select **Start at** and specify the date and time parameters.

3. Click **OK** to save the settings.

Application connection settings

You can configure routing and mirroring settings for your application connection. In addition, there are application specific settings that you can configure. The fields on the **Connection** tab will vary depending on the type of application you are protecting.

- Routing data transmissions
- Protection configuration
  - Configuring Exchange storage group protection
  - Configuring SQL database protection
  - Configuring file share protection
  - Configuring BlackBerry database protection
  - Configuring SharePoint database protection
- Mirroring data

Routing data transmissions

1. Make sure you have a valid domain and servers specified, click **Configure** from the main Application Manager page, and then select the **Connection** tab.

   Note: The fields on the **Connection** tab will vary depending on the type of application you are protecting.

2. By default, Double-Take Availability will select the default **Route** for transmissions. If desired, select a different IP address on the target that will be used for transmissions. If you are using a cluster, use the virtual server IP address.

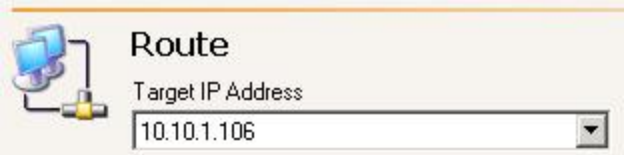

3. Click **OK** to save the settings.

Protection configuration

Application specific protection options are available on the **Configure Protection Connection** tab. The fields on the **Connection** tab will vary depending on the type of application you are protecting.

- Configuring Exchange storage group protection
- Configuring SQL database protection
- Configuring file share protection
- Configuring BlackBerry database protection
- Configuring SharePoint database protection

Configuring Exchange storage group protection

If you are protecting Exchange, you can specify which storage groups, mailboxes, and public folder stores that you want to protect.

1. Make sure you have a valid domain and servers specified, click **Configure** from the main Application Manager page, and then select the **Connection** tab.

Note: The fields on the **Connection** tab will vary depending on the type of application you are protecting.

---

2. The **Protected Storage Groups** will list the storage groups, mailboxes, and public folder stores. By default everything is slected. Select the storage groups that you want to protect. By selecting individual storage groups, you can reduce the amount of data being replicated and filter out storage groups that do not need to be protected or failed over. Only the users associated with the selected storage groups will be failed over.

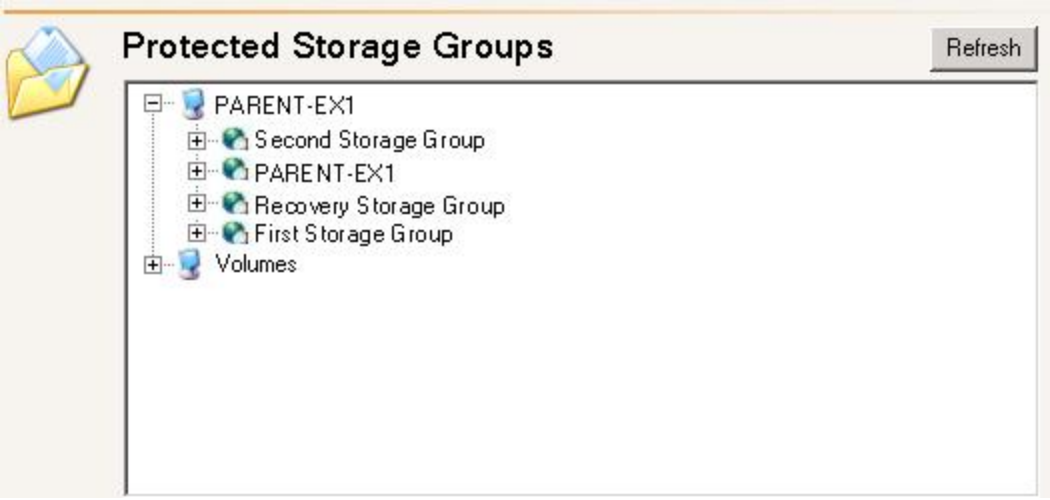

Note: If you do not select all storage groups, you should make sure that other backups are available.

Ideally, you should place all query-based distribution groups in a single organization container and give the target server full control to the container and all child objects.

Names with a plus sign (+) are not supported. You must rename the storage group and remove the plus sign.

The **Protected Storage Groups** list will be disabled if you have enabled **Override Generated Rules** on the Advanced tab.

3. If desired, you can select additional data to protect under the **Volumes** folder.
4. Click **Refresh** if you need to refresh the items in the tree view.
5. Click **OK** to save the settings.

Configuring SQL database protection

If you are protecting SQL, you can protect the SQL instance or only the database only.

1. Make sure you have a valid domain and servers specified, click **Configure** from the main Application Manager page, and then select the **Connection** tab.

Note: The fields on the **Connection** tab will vary depending on the type of application you are protecting.

---

2. Select if you want to protect the **SQL Instance** or the **Database Only**. The **Protected Databases** options and list will be disabled if you have enabled **Override Generated Rules** on the Advanced tab.

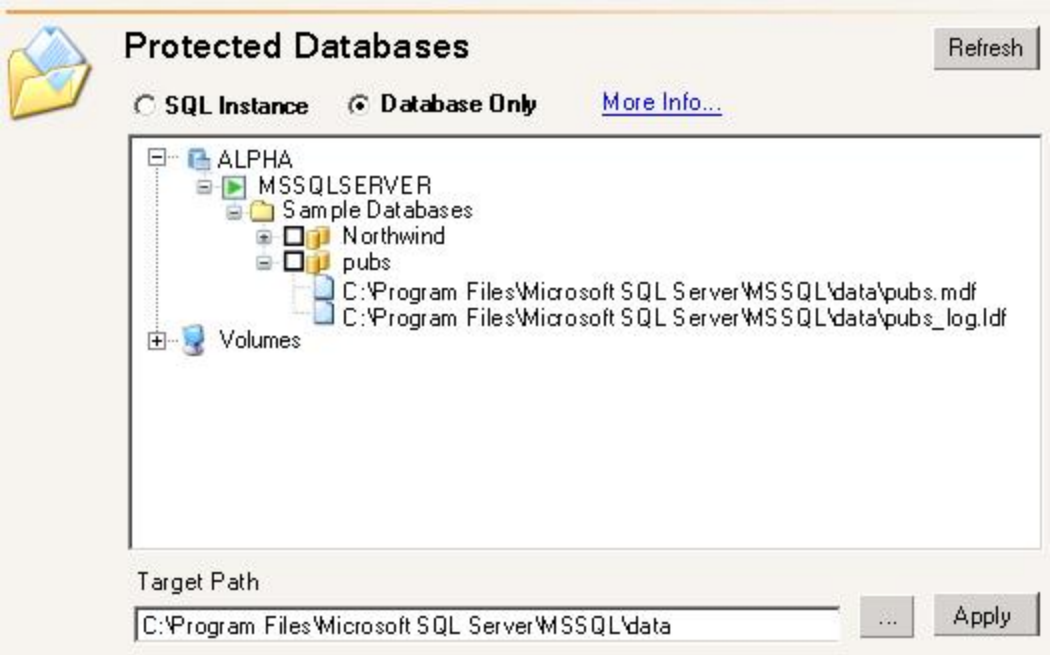

- **SQL Instance**—This option will protect the entire SQL program and data files (except the \binn directory). With this option, end-users can access the SQL data from the target in the event of a failure. With this option, the source and target servers must have the following configuration.
  - The servers must have the same version of SQL (major and minor versions).
  - The servers must have the same logical drive structure where the SQL program and data files are stored.

- The servers must have the same named instances, unless you are running Application Manager in advanced mode. In this case, you can identify instances to protect that are offline or do not exist on the target. The Manage SQL Server Instances dialog will only appear if you two or more SQL instances (default plus one or more named instances or two or more named instances with no default instance).
- The TcpPort for the named instances will be different. This is acceptable.
- You can exclude user databases from protection, but the system databases (except for tempdb) are required.
- You may want to exclude the tempdb database to reduce mirroring and replication traffic.

- **Database Only**—This option will protect the .mdf, .ldf, and .ndf files. With this option, the databases will be attached to the target in the event of a failure and then end-users an access the data. This option is intended for advanced users only. During the validation process, you will have the opportunity to transfer user logins and permissions (both server and database-level) and certain SQL Server registry and configuration settings to the target server. This will allow users to access the data associated with the selected database(s), but no other server-level functionality will be transferred to the target server, including but not limited to Job Server configuration, Full-Text service configuration, SQL Replication configuration, linked servers, remote servers, and backup devices.

  With this option, the source and target servers must have the following configuration.

  - You must configured any SQL Server replication on the protected source databases on the target after failover.
  - Transparent Data Encryption (TDE) is not supported for SQL 2008 when using database only mode.
  - Attempting to attach a replicated SQL database on the target server after failover outside of the Application Manager can fail.
  - The Double-Take service account (typically the target's LocalSystem account) is the account used to attach and detach databases on failover and failback. When the database is detached by the failover and failback scripts, the Double-Take service account becomes the owner of those files that make up the database (*.mdf, *.ldf, and so on). Any attempts to manually attach the database may fail if the user account does not yet have NTFS permissions to access the physical files. To change the permissions on an individual file, perform these steps on each file that is part of the database's file list.

a. In Windows Explorer, right-click the folder that contains the physical files for the database that needs to be manually attached.
b. Select **Properties**.
c. On the **Security** tab, determine if the user account has NTFS permissions for that folder.
d. If the user account does not have specific or inherited permissions, click the **Add** button.
e. Enter the user account name (such as domain\administrator).
f. After the user account has been added, give the account **Full Control** permissions.
g. Make sure that the subfolders and files are set to inherit these rights, then click **OK**.

- If you are working in a many-to-one scenario, and you have two SQL servers and each has only the default instance installed, you can protect databases from both servers' default instances provided that the database names are unique. For example, if both servers' default instances have a database named Accounting, you can only protect and failover one server's copy of the database because SQL on the target will not allow you to attach more than one copy of the same-named database. The first server to failover will attach its Accounting database, while the second server to failover will not attach its Accounting database.

  If your two SQL servers have unique instances installed, you can protect databases from both servers if the target has at least those two instances installed.

  Keep in mind, if the database names (accounting1.mdf and accounting2.mdf) or locations on the target (\source1\accounting1\accounting.mdf and \source2\accounting2\accounting.mdf) are unique, you can protect and failover both databases to the same target.

3. If you selected **Database Only**, you can highlight a non-system database and then identify a unique **Target Path**. Keep in mind that if you are protecting multiple databases, specifying a unique target path will impact the location of any databases in the same or lower directory structure. For example, if you specify a unique target path for \level1\level2\database.mdf, the target path for level1\level2\level3\database.mdf will also be in that path. Click the ellipse (**...**) button to search for a target path and click **Apply** to save the setting.
4. If desired, you can select additional data to protect under the **Volumes** folder.

5. Click **Refresh** if you need to refresh the items in the tree view.
6. Click **OK** to save the settings.

Configuring file server protection

If you are protecting a file server, you can specify the file shares that you want to protect.

1. Make sure you have a valid domain and servers specified, click **Configure** from the main Application Manager page, and then select the **Connection** tab.

Note: The fields on the **Connection** tab will vary depending on the type of application you are protecting.

---

2. By default, all non-administrative file shares will be selected. Select or deselect the file shares in the tree that you want to protect.

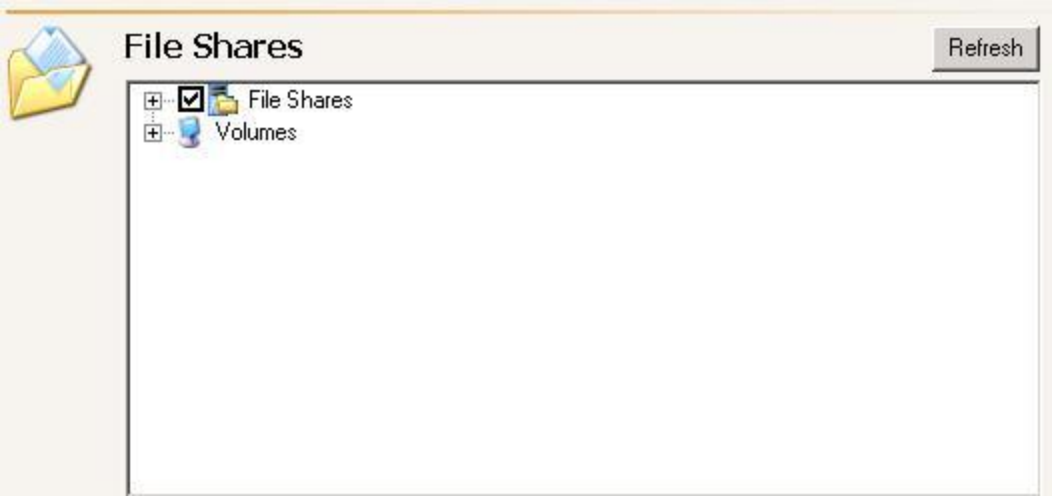

Note: The **File Shares** list will be disabled if you have enabled **Override Generated Rules** on the Advanced tab.

If your source is a domain controller, you cannot protect the NETLOGON and SYSVOL shares and they will not be visible in the File Shares tree.

---

3. If desired, you can select additional data to protect under the **Volumes** folder.
4. Click **Refresh** if you need to refresh the items in the tree view.
5. Click **OK** to save the settings.

Configuring BlackBerry database protection

1. Make sure you have a valid domain and servers specified, click **Configure** from the main Application Manager page, and then select the **Connection** tab.

Note: The fields on the **Connection** tab will vary depending on the type of application you are protecting.

---

2. Select the databases that you want to protect.

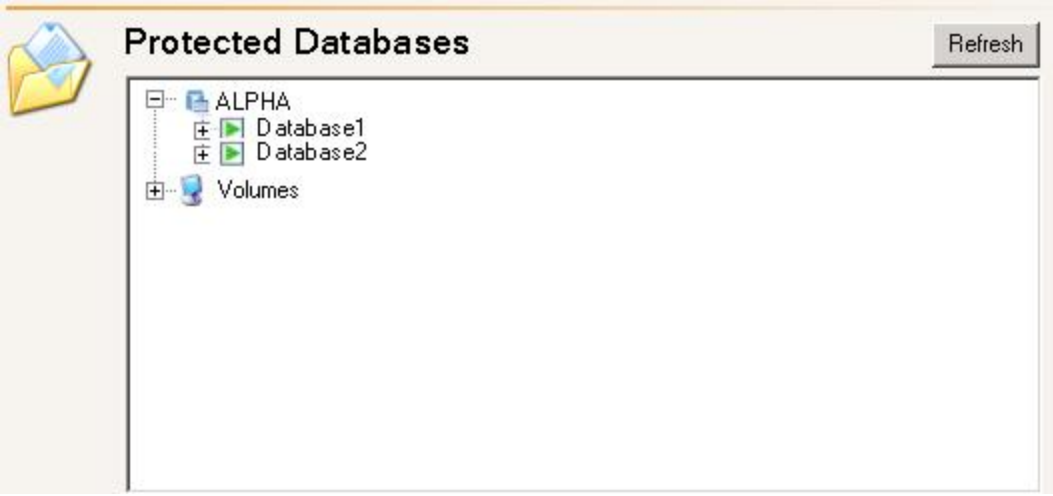

Note: You cannot deselect the databases containing the BES or MDS information.

If you do not want to protect the MDS database, deselect **Protect MDS services** on the **BlackBerry** tab using the instructions below.

You may want to exlude the tempdb database to reduce mirroring and replication traffic.

The **Protected Databases** list will be disabled if you have enabled **Override Generated Rules** on the Advanced tab.

---

3. If desired, you can select additional data to protect under the **Volumes** folder.
4. Click **Refresh** if you need to refresh the items in the tree view.
5. Next select the **BlackBerry** tab.
6. Under **BlackBerry Options** select to **Protect log files** and/or **Protect MDS**

service(s).

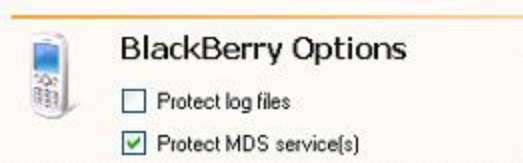

7. Under **BlackBerry Services**, Application Manager automatically determines the appropriate BlackBerry services to start and stop. If necessary, modify the list of services.

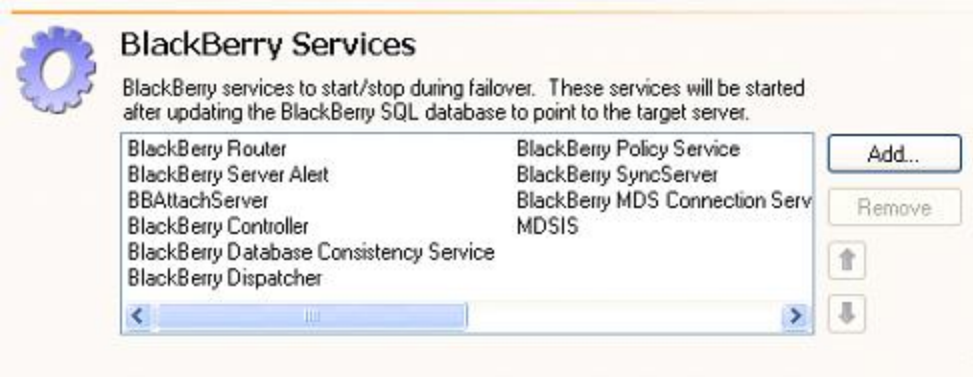

   a. Click **Add** to insert a service into the list. Specify the service name and make sure that **Service must be stopped on target** is enabled so that replication can update files on the target.
   b. Click **Remove** to remove a service from the list. You can only remove services that you have manually added.
   c. If you are configuring your services, highlight a service and click the up or down arrow to reorder the list. The services will be stopped and started in the order displayed.
8. Click **OK** to save the settings.

Configuring SharePoint database protection

1. Make sure you have a valid domain and servers specified, click **Configure** from the main Application Manager page, and then select the **Connection** tab.

Note: The fields on the **Connection** tab will vary depending on the type of application you are protecting.

---

2. By default, the entire SharePoint program and data files (except the \binn directory) will be protected. End-users can access the data from the target in the event of a failure.

Note: The servers must have the same version of SharePoint (major and minor versions).

The servers must have the same logical drive structure where the SharePoint program and data files are stored.

If you are using a SQL server named instance for a back-end database server in your SharePoint setup, both the source and target SQL servers must have the same named instances and the same logical drive structure.

You may want to exlude the tempdb database to reduce mirroring and replication traffic.

The **Protected Databases** list will be disabled if you have enabled **Override Generated Rules** on the Advanced tab.

3. If desired, you can select additional data to protect under the **Volumes** folder.
4. Click **Refresh** if you need to refresh the items in the tree view.
5. Next select the **SharePoint** tab. These options allow you to join or extend the target front-end web server to the production SharePoint configuration or web farm. Application Manager determines the Microsoft SQL server and configuration database used by the source SharePoint web front-end server, then uses that information to connect the specified target web server to the same SharePoint configuration. The target web server specified can be local or remote.

Note: The target web server must have the same version of SharePoint installed as the product SharePoint web server.

For best results, SharePoint should be installed but not yet configured on the target web server.

In order to extend the target web server, you will need to add the SharePoint administrator account to the local domain administrator group on the target server before you extend target web front-end server into the farm.

6. The first five fields will be filled in automatically. However, you can modify any of the fields.
   - **Server Name**—Enter the NetBIOS or physical name of the target SharePoint web server.
   - **IP Address**—Enter the IP address for the target web server.
   - **TCP Port**—Enter the TCP port to be used for communicating with the target web server.
   - **Config Database Server**—Enter the name of the Microsoft SQL Server that hosts the configuration database.
   - **Config Database Name**—Enter the name of the configuration database for the production SharePoint web front-end server.
   - **SharePoint Admin Name**—Enter the account used to install and configure SharePoint on the production SharePoint web front-end server. This should be entered as a fully qualified domain name in the format domain\username.

- **SharePoint Admin Password**—Enter the password for the SharePoint Admin account.
- **Confirm Password**—Re-enter the password.

7. After you have entered the parameters, click **Connect Server** and the specified SharePoint front-end web will be extended into the source SharePoint configuration. This process may take several minutes to complete. During this time, you can perform other tasks within the **Configure Protection** window, however, you will not be able to close the **Configure Protection** window until the task is complete.

Note: You must manually install the Central Administration web application after the target has been extended in order to be able to administer SharePoint in the event of a failover.

---

8. Click **OK** to save the settings.

Mirroring data

1. Make sure you have a valid domain and servers specified, click **Configure** from the main Application Manager page, and then select the **Connection** tab.

Note: The fields on the **Connection** tab will vary depending on the type of application you are protecting.

---

2. Specify your **Mirror Settings**.

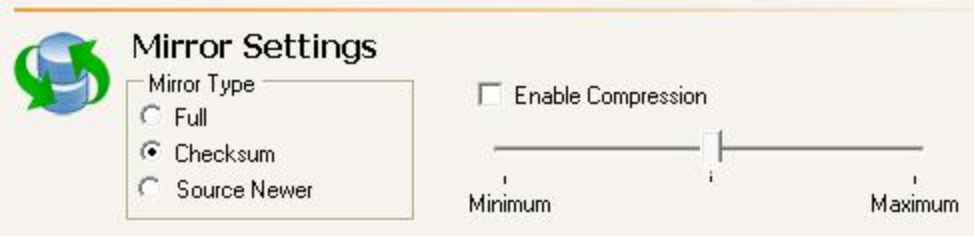

- **Mirror Type**—Select the type of Double-Take Availability mirroring process you want to perform. A **Full** mirror will transmit all files from the source to the target. A **Checksum** mirror will transmit only the blocks of data that are different between the source and target. A **Source Newer** mirror will transmit only those files that are newer on the source than on the target. The newer option is only available for file server protection when you have launched the Application Manager using the command line advanced option.
- **Enable Compression**—Compression allows you to reduce the amount of bandwidth needed to transmit data from the source to the target. The data is compressed before being transmitted and then is uncompressed before it is written on the target. Typically, compression is used in WAN environments, but not in LAN environments. If desired, enable compression and select the level of compression that you want to use. All connections to the same target will have the same compression settings.

3. Click **OK** to save the settings.

Application advanced settings

You can configure advanced settings for you application connection. The advanced settings that are available will depend on the application you are protecting. Therefore, the fields on the **Advanced** tab will vary. In addition, the fields will vary depending on if you launched Application Manager in standard or advanced mode.

- Configuring the replication set
- Configuring scripts
- Configuring Active Directory
- Configuring items to failover
- Configuring default connection parameters

Configuring the replication set

Application Manager automatically creates a replication set with a name based on the application you are protecting. The list below contains the default replication set names, where source and target are the names of the respective servers.

- **Exchange**—xdag01_source_target
- **SQL**—sqldag01_source_target
- **File server**—fileprint_source_target
- **BlackBerry**—BB_source_target
- **SharePoint**—SharePointdag01_source_target

Application Manager selects all of the necessary directories and files to add to your replication set to protect your application. You should only modify the replication set definition if there are additional directories or files that you want to protect. Do not modify the rules unless you are familiar with Double-Take Availability and your application.

Exchange Note: If you are protecting Exchange and want to protect the Badmail folder, you will need to manually add it using the instructions below.

SQL Note: If you are protecting SQL, the folder that contains a database's FILESTREAM data will automatically be included in the replication set for protection if the database is included for protection. Any steps previously taken to enable FILESTREAM support on the source (for instance, file system or operating system-level changes that were made specifically to support FILESTREAM) must also be applied similarly to the target for consistency. Failure to account for FILESTREAM-specific changes on the target server, or specifically excluding FILESTREAM data files/paths from the replication set, can impact SQL Server's ability to mount a database with FILESTREAM data when failing over.

---

1. Make sure you have a valid domain and servers specified, click **Configure** from the main Application Manager page, and then select the **Advanced** tab.

Note: The fields on the **Advanced** tab will vary depending on the type of application you are protecting. In addition, the fields will vary depending on if you launched Application Manager in standard or advanced mode.

2. To modify the replication set definition, select **Override Generated Rules**. When this option is selected, the application controls (storage groups or protected databases) on the **Connection** tab will be disabled.
3. To add a new replication set rule, click **Add**.
4. Specify a path, wild card, or specific file name. Application Manager will not verify the rule you are adding.
5. Select to **Include** or **Exclude** the file.
6. Mark the rule as **Recursive** or **Non-Recursive** if you want the rule applied to subdirectories.
7. Click **Add**.
8. Repeat steps 4 through 7 for each replication set rule you want to add.
9. When you are finished adding rules, click **Close**.
10. If you need to remove a rule, highlight it and click **Remove**. If you remove a rule added by Application Manager, you could impact the success of failover.
11. Click **OK** to save the settings.

Configuring scripts

Scripts are executed at different points during failover, failback, and restoration. The scripts perform actions to make your applications available on the appropriate server. Editing scripts is an advanced feature. Do not modify the scripts unless you are familiar with Double-Take Availability, your application, and scripting. Any edits should be made carefully and tested prior to deployment to ensure the changes are correct. Incorrect script changes could cause failover issues.

1. Make sure you have a valid domain and servers specified, click **Configure** from the main Application Manager page, and then select the **Advanced** tab.

   Note: The fields on the **Advanced** tab will vary depending on the type of application you are protecting. In addition, the fields will vary depending on if you launched Application Manager in standard or advanced mode.

   ---

2. Click on the button associated with the script you want to edit.
   - **Failover Script**—This script is executed automatically on the target after the core failover processes have completed.
   - **Failback Script**—This script is executed automatically on the target before the failback processes begin.
   - **Restore Script**—This script is not executed automatically, but is available for the source if needed.
   - **Post Failback Script**—This script is not executed automatically, but is available for the target if needed.

     Any changes you save to the scripts will be copied to the appropriate server when the configuration changes are accepted. If you reconfigure your application protection after making script changes, Application Manager will copy updated scripts to the appropriate server, overwriting any changes that you manually made. You should make a backup copy of your script changes to copy over after making Application Manager updates. If you want to make the script changes permanent, you must modify the script files manually in the Double-Take Availability installation location.
3. Click **OK** to save the settings.

Configuring Active Directory

If you are protecting Exchange, you can configure several settings for Active Directory.

1. Make sure you have a valid domain and servers specified, click **Configure** from the main Application Manager page, and then select the **Advanced** tab.

Note: The fields on the **Advanced** tab will vary depending on the type of application you are protecting. In addition, the fields will vary depending on if you launched Application Manager in standard or advanced mode.

---

2. Enable **Force AD replication** if you want Active Directory replication to be initiated from the source and target's domain controller. Each time the Double-Take Availability Exchange Failover utility is executed in the failover and failback scripts, Active Directory replication will be forced. If you do not want to force Active Directory replication, disable this option.
3. If you have enabled forced Active Directory replication, specify the **Max wait time for AD replication**. This is the length of time, in minutes, that failover or failback will wait for before continuing. If replication exceeds the wait time specified, a log created and replication and failover continue. This wait time does impact when failover and failback will complete, however it does not impact the success or failure.
4. If you are using Exchange 2003, you can specify the name (not an IP address) of the **Target domain controller** which is the server where updates will be made during failover and failback. If you do not specify a domain controller, then the domain controller determined by Active Directory will be used.
5. Click **OK** to save the settings.

Note: If you want to add the target back to the PF list to which the source belongs, you will need to enable the **Restore PF Tree** option.

   a. From the main Application Manager screen, select **Tools**, **Actions**.
   b. Enable **Display Advanced Options**.
   c. Reselect **Protect Exchange Server** from the **Tasks** list in the left pane and the **Restore PF Tree** option will be added to the **Actions** menu.
   d. Select **Actions**, **Restore PF Tree**. This will copy the owning PF tree setting from the source public folders to the target public folders.

This setting is cleared when protection is enabled, which prevents SMTP queuing issues when trying to deliver messages to the target, but is never restored. If you want to have an active target server, you can use this command to restore it to an Application Manager state.

Configuring items to failover

If you are performing an identity failover, then you have already selected the **Items to Failover**. Changing any of the **Items to Failover** on the **Advanced** tab will automatically make the same change on the **Failover** tab **Configure Identity Failover** page. However, if you are performing DNS failover, you may want to modify the items that you are failing over using the instructions below.

1. Make sure you have a valid domain and servers specified, click **Configure** from the main Application Manager page, and then select the **Advanced** tab.

   Note: The fields on the **Advanced** tab will vary depending on the type of application you are protecting. In addition, the fields will vary depending on if you launched Application Manager in standard or advanced mode.

2. The default selected items under **Items to Failover** will depend on the application you are protecting. Enable or disable if you want to failover the source's server name, file shares, and/or Active Directory host name.

   Note: If your source and target are on different subnets, you should not failover the IP address.

   Exchange Note: Do not failover the Active Directory host name.

   SQL Note: Do not failover the server name.

   File Server Note: You cannot failover file shares from a parent to child domain.

3. Click **OK** to save the settings.

Configuring default connection parameters

Several default connection options are available which allow you to disable or enable the creation of default connection parameters. Ideally, you want Application Manager to create the default parameters. However, if you have modified any of the parameters manually and you do not want your modifications overwritten by the defaults, then you may want to disable the creation of the default connection parameters.

1. Make sure you have a valid domain and servers specified, click **Configure** from the main Application Manager page, and then select the **Advanced** tab.

Note: The fields on the **Advanced** tab will vary depending on the type of application you are protecting. In addition, the fields will vary depending on if you launched Application Manager in standard or advanced mode.

---

2. **Create Replication Set** indicates if the default replication set generated by Application Manager will be used to establish an application protection connection. If enabled, the default replication set will be generated and used. Disable this option if you have manually updated the default replication set.
3. **Create Failover Scripts** indicates if the default failover, failback, and restore scripts generated by Application Manager will be used during those processes. If enabled, the default scripts will be used. Disable this option if you have manually updated the scripts.
4. **Create Connection** indicates if a connection will be established when application protection is enabled. Disable this option if you want to perform testing on your replication set prior to establishing the connection.
5. **Create Failover Monitor** indicates if failover monitoring will be started when application protection is enabled. Disable this option if you want to perform testing on your scripts prior to establishing monitoring.
6. Click **OK** to save the settings.

Using NAT or firewalls with application workloads

If your source and target are on opposite sides of a NAT or firewall, you will need to configure your hardware to accommodate application workload communications. You must have the hardware already in place and know how to configure the hardware ports. If you do not, see the reference manual for your hardware.

- Application workload ports
- Microsoft Windows ports
- Hardware ports

Application workload ports

By default, Double-Take Availability uses port 6320 for all communications. To verify or modify the ports, use the following instructions.

1. In the Application Manager, select **Tools**, **Options**.
2. Verify or modify the **Service Listen Port** as needed. All servers and clients in your application workload protection must be using the same port.

Double-Take Availability uses ICMP pings to monitor the source for failover. A failover monitor will not be created if ICMP is blocked (although the data and application will still be protected). You should configure your hardware to allow ICMP pings between the source and target. If you cannot, you will have to monitor for a failure using the Double-Take Availability replication service. Use the **Method to Monitor for Failover** option on the configuration **Monitoring** tab to set your failover monitoring method.

SharePoint Note: You also need to configure your hardware to allow communication on port 6350.

Microsoft Windows ports

Application workload protection uses WMI (Windows Management Instrumentation) which uses RPC (Remote Procedure Call). By default, RPC will use ports at random above 1024, and these ports must be open on your firewall. RPC ports can be configured to a specific range by specific registry changes and a reboot. See the Microsoft Knowledge Base article 154596 for instructions.

Application workload protections also rely on other Microsoft Windows ports.

- Microsoft File Share uses ports 135 through 139 for TCP and UDP communications.
- Microsoft Directory uses port 445 for TCP and UDP communications.

These ports must be open on your firewall. Check your Microsoft documentation if you need to modify these ports.

Hardware ports

You need to configure your hardware so that the application workload ports and Microsoft Windows ports are open. Since communication occurs bidirectionally, make sure you configure both incoming and outgoing traffic.

There are many types of hardware on the market, and each can be configured differently. See your hardware reference manual for instructions on setting up your particular router.

Exchange Failover Utility

When you configure Exchange for protection, Application Manager creates customized scripts based on the settings you choose. The scripts are based on the Exchange Failover Utility (EFO). Generally, you do not need to run the Exchange Failover Utility from the command line or modify the scripts that Application Manager creates, however, the syntax for the utility is provided below in case that need arises.

Command

EFO

Description

Used in script files to failover Exchange data

Syntax

EXCHFAILOVER -FAILOVER | -FAILBACK -S <***source***> -T <***target***> [-L <***filename***>] [-NORUS] [-NORM] [-NOSPN] [-NOOAB] [-NOADREPLICATION] [-MAXREPWAIT <***minutes***>] [-NOEXCHANGEAB] [-NOQUERYBASEDDISTGROUPS] [-NORGCONNECTORS] [-NOPUBLICFOLDERS] [-ONLYPUBLICFOLDERS] [-MOVEHOSTSPN] [-O <***filename***>] [-R [<***source_group***>] [,<***source_mailstore***>] [: [<***target_group***>] [,<***target_mailstore***>]]] [-SETUP] [-TEST] [-U <***name***> : <***password***>] [-VIRTUAL <***new_IPaddress***>] [-DC <***domain_name***> | <***IPaddress***>] [-SDOMAIN] [-TDOMAIN] [/?] [/??]

Options

- FAILOVER—Exchange data will be moved from the source to the target
- FAILBACK—Exchange data will be moved from the target to the source. Even through the flow of data has changed (target to source), the source-related options with this utility still pertain to your original source (or a new source if you had to replace the source). The target-related options pertain to the original target, the server that is currently standing in for the source.
- S ***source***—Name of the original source
- T ***target***—Name of the original target
- L ***filename***—Name of the log file. 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
- NORM—Do not change the Routing Master
- NOSPN—Do not change the Service Principal Name
- NOOAB—Do not change the siteFolderServer for the offline address book
- NOADREPLICATION—Do not force Active Directory replication
- MAXREPWAIT ***minutes***—Maximum time, in minutes, to wait for Active Directory replication to complete before continuing failover or failback. The default value is 30 minutes. This option is not applicable if NOADREPLICATION is used.
- NOEXCHANGEAB—Do not fail back the ExchangeAB Service Principal Name for Small Business Server
- NOQUERYBASEDDISTGROUPS—Do not update query-based distribution lists
- NORGCONNECTORS—Do not change Routing Group connectors
- NOPUBLICFOLDERS—Do not move public folders
- ONLYPUBLICFOLDERS—Only move public folders
- MOVEHOSTSPN—Move the HOST Service Principal Name to or from the target instead of removing and adding it
- O ***filename***—Name of the file that contains the options to pass through to the Exchange Failover utility
- R ***source_group***, ***source_mailstore*** : ***target_group*** , ***target_mailstore***—By itself, this option creates a one-to-one mapping of the groups and mail stores from the source to the target. Optionally, you can supply group and mail store names to customize the mapping. Repeat this option as often as needed.
- SETUP— Sets the **overwrite database on restore** flag without completing user moves or RUS and folder updates. If this option is not used, the Exchange Failover utility still sets the **overwrite database on restore** flag, but the other work is also performed.
- TEST—Runs in test mode, so no Active Directory updates are made
- U ***name*** : ***password***—User with Active Directory permissions. The password is case-sensitive.
- VIRTUAL ***new_IPaddress***—Performs updates to virtual protocols only for like-named cluster failover

- DC *domain_name* | *IPaddress*—Name or IP address of the domain controller to make updates on
- SDOMAIN—Source domain name
- TDOMAIN—Target domain name
- ?—Displays the Exchange Failover utility syntax
- ??—Displays a description of the Exchange Failover utility options

Examples

- exchfailover -failover -s alpha -t beta
- exchfailover -failback -s alpha -t beta
- exchfailover -failover -s alpha -t beta -r -onlypublicfolders
- exchfailover -failover -s alpha -t beta -u administrator:password -dc domaincontroller

---

Virtual server protection

Select a link to jump to instructions that correspond with your virtual configuration.

- **Physical or virtual (guest level) to Hyper-V or ESX virtual**—If your source is a physical or virtual server, and you want to protect the volumes from the physical server or the volumes from within the virtual guest operating system, and your target is a virtual server on a Hyper- V or ESX server, then see Protecting an entire physical or virtual server to a Hyper-V or ESX server.
- **Hyper-V virtual (host level) to Hyper-V virtual**—If your source is a virtual server on a Hyper-V server, and you want to protect the host-level virtual disk files (.vhd files), and your target is a virtual server on a Hyper-V server, then see Protecting a Hyper-V server to a Hyper-V server.
- **ESX virtual (host level) to ESX virtual**—If your source is a virtual server on an ESX server, and you want to protect the host-level virtual disk files (.vmdk files), and your target is a virtual server on an ESX server, then see Protecting an ESX server to an ESX server.

Protecting a physical or virtual server to a Hyper-V or ESX server

Use these instructions to protect a physical or virtual server, where you want to protect the volumes from the physical server or the volumes from within the virtual guest operating system, to a virtual server on a Hyper-V or ESX server.

1. Open the Double-Take Console by selecting **Start**, **Programs**, **Double-Take**, **Double-Take Console**.
2. Click **Get Started** from the toolbar.
3. Select **Double-Take Availability** and click **Next**.
4. Select **Protect an entire server using a Hyper-V or ESX virtual machine** and click **Next**.
5. Specify the source server that you want to protect. This is the physical or virtual server that you want to protect.

- **Server**—Specify the name or IP address of the physical or virtual server that you want to protect. You can click **Browse** to select a server from a network drill-down list.

Note: If you enter the source server's fully-qualified domain name, the Double-Take Console will resolve the entry to the server short name. If that short name resides in two different domains, this could result in name resolution issues. In this case, enter the IP address of the server.

Your source can have no more than four NICs enabled, except if your source is ESX version 4.x. In that case, you can have up to ten NICs.

- **User name**—Specify a user that is a member of the Double-Take Admin and local administrator security groups on the source. If you want to use a domain user account, enter a fully-qualified domain name in the format domain\username or username@domain.
- **Password**—Specify the password associated with the **User name** you entered.

6. Click **Next** to continue.
7. Choose the volumes on the source server that you want to protect.

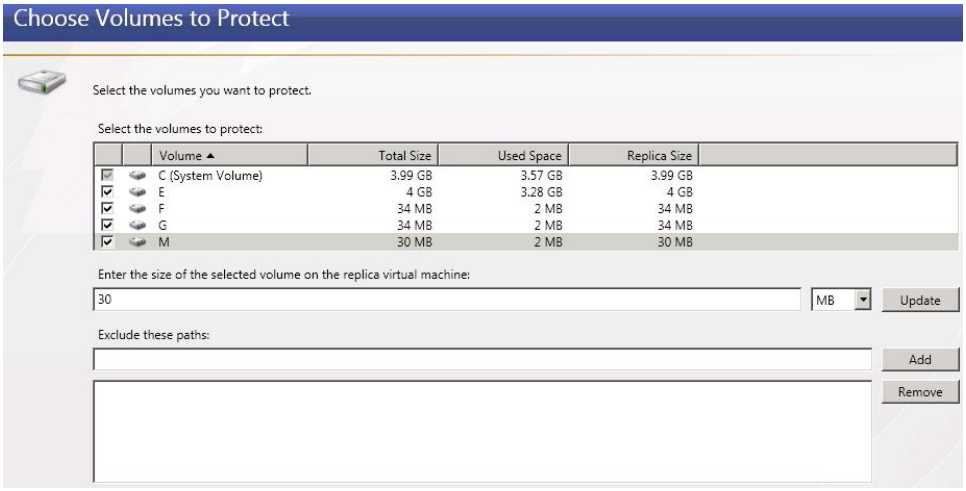

- **Select the volumes to protect**—By default the system volume will be selected for protection. You will be unable to deselect the system volume. Select any other volumes on the source that you want to protect.
- **Enter the size of the selected volume on the replica virtual machine**—Highlight a volume that you are protecting and specify the size, including the value in MB or GB, of the replica virtual machine on the target. The value must be at least the size of the specified **Used Space** on that volume. Click **Update** to refresh the **Replica Size** value in the table. Repeat this process for each volume you are protecting.

Note: If the size of the replica virtual machine is identical to the size of the source volume and the source has less than 20 MB of free disk space remaining, you may run out of disk space on the replica due to differences in how the virtual disk's block size is formatted. To avoid

this issue, specify the size of your replica virtual machine to be at least 20 MB larger.

- **Exclude these paths**—If there are specific paths on a volume that you do not want to protect, specify those locations. Keep in mind that missing data may impact the integrity of your applications.

8. Click **Next** to continue.
9. Specify the target server. If you are protecting to a Hyper-V server, this is the name of the Hyper-V server. If you are protecting to a ESX server, this is the host name of your virtual recovery appliance.

- **Server**—Specify the name or IP address of the target server. You can click **Browse** to select a server from a network drill-down list.

Note: If you enter the target server's fully-qualified domain name, the Double-Take Console will resolve the entry to the server short name. If that short name resides in two different domains, this could result in name resolution issues. In this case, enter the IP address of the server.

- **User name**—Specify a user that is a member of the Double-Take Admin security group. If your target is a Hyper-V server, the user must also have administrative rights for Microsoft Hyper-V. If you want to use a domain user account, enter a fully-qualified domain name in the format domain\username or username@domain.

- **Password**—Specify the password associated with the **User name** you entered.

10. Click **Next** to continue.
11. If you are protecting to an ESX server, specify the ESX server information for your target virtual recovery appliance.

- **VirtualCenter Server**—Select your VirtualCenter server from the list. If your VirtualCenter server is not in the list, click **Add VirtualCenter Server**, specify the server and valid credentials, and click **Add**. If you are not using VirtualCenter, select **None**.
- **ESX Server**—Specify the name or IP address of the ESX server.
- **User name**—This field will only be available if you are not using VirtualCenter. In this case, specify the root user or another user that has the administrator role on the specified ESX server.
- **Password**—Specify the password associated with the **User name** you entered.

12. Click **Next** to continue.
13. If you are using a Hyper-V target, select a location on the target for the replica virtual machine.

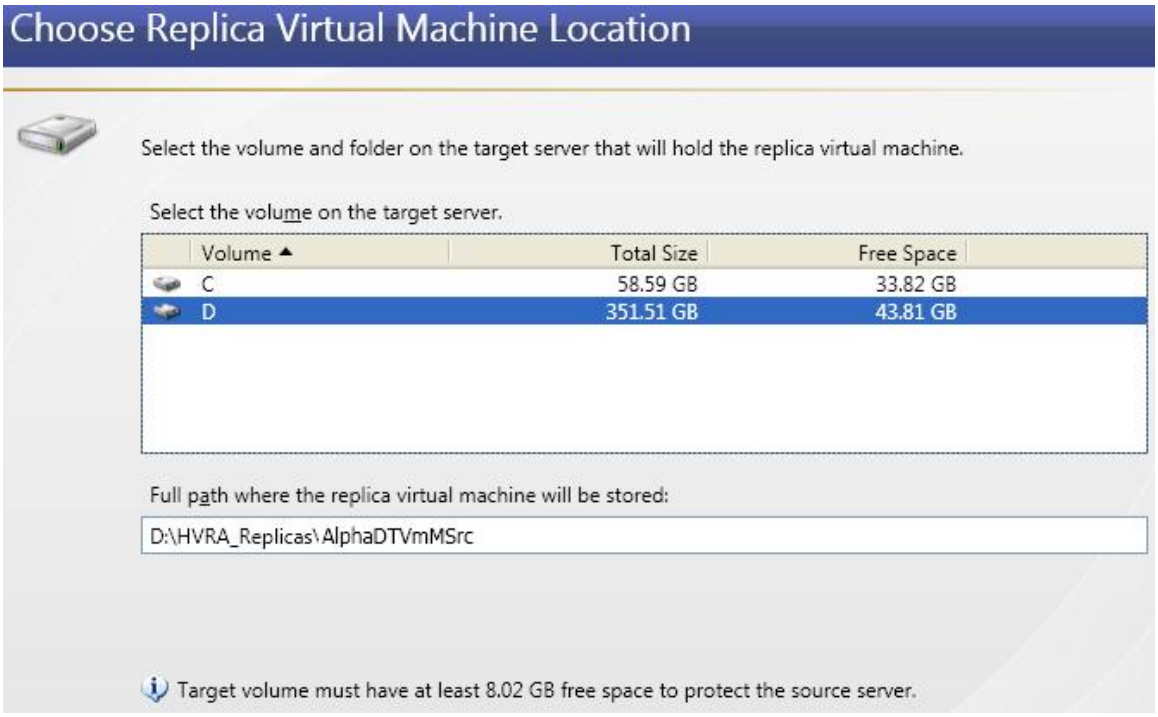

- **Volume**—Select one of the volumes from the list to indicate which volume on the target where you want to store the new virtual server when it is created. The target volume must have enough **Free Space** to store the source data. The minimum size is noted at the bottom of the page.
- **Full path where the replica virtual machine will be stored**—Specify a location on the selected **Volume** to store the replica of the source. By specifying an existing folder, you can reuse an existing virtual machine on your Hyper-V target created by a previous protection job. This can be useful for pre-staging data on a virtual machine over a LAN connection and then relocating it to a remote site after the initial mirror is complete. When you reuse a virtual machine, Double-Take Availability performs a difference mirror which saves time. Use the following steps to reuse a virtual machine on a Hyper-V target.
    a. Create a protection job in a LAN environment, letting Double-Take Availability create the virtual disk for you.
    b. Complete the mirror process locally.
    c. Disconnect and delete the protection job and when prompted, select to **Keep and register the replica virtual machine**.
    d. From the Hyper-V manager, delete the replica virtual machine, which will delete the virtual machine configuration but will keep the associated .vhd files.
    e. Shut down and move the Hyper-V target server to your remote site.

f. After the target server is back online at the remote site, create a new protection job for the same source server and select to reuse the existing virtual machine when prompted. Double-Take Availability will reuse the existing .vhd files and perform a difference mirror over the WAN to bring the virtual machine up-to-date.

14. If you are using an ESX target, select a location on the target for the replica virtual machine.

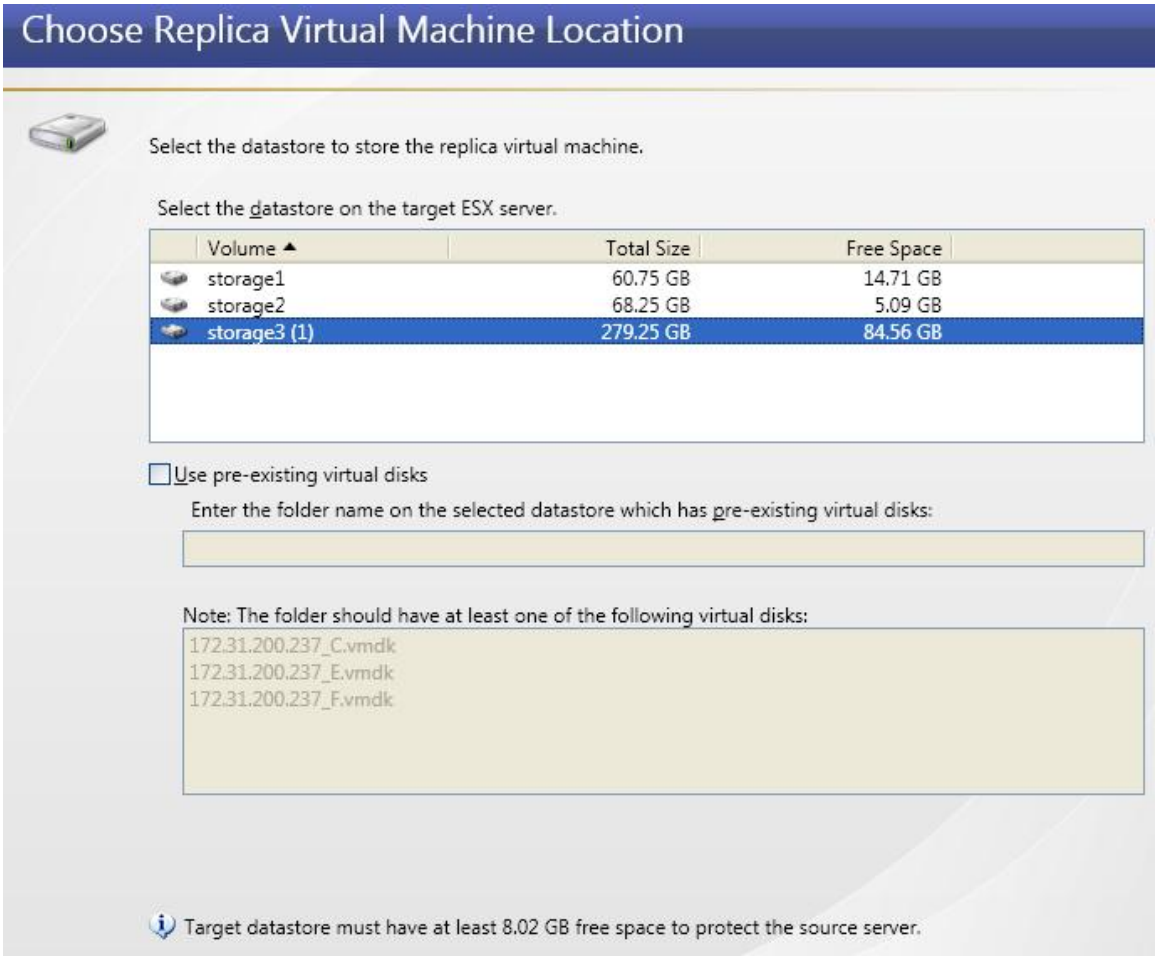

- **Volume**—Select one of the volumes from the list to indicate which volume on the target where you want to store the new virtual server when it is created. The target volume must have enough **Free Space** to store the source data. The minimum size is noted at the bottom of the page.
- **Use pre-existing virtual disks**—You can reuse an existing virtual disk on your ESX target, rather than having Double-Take Availabilitycreate a virtual disk for you. This saves time by skipping the virtual disk creation steps and performing a difference mirror instead of a full mirror. In order to use a pre-existing virtual disk, it must be a valid VMware virtual disk. It cannot be attached to any other virtual machine, and the virtual disk size cannot be

changed.

Because Double-Take Availability will skip the virtual disk creation steps when using a pre-existing disk, Double-Take Availability will instead copy your existing virtual disk to the default VMware new virtual machine location. Therefore, it is important that you do not place your existing virtual disk in the new folder location that VMware will create. Put the pre-existing virtual disk in a temporary location on the target. Specify this temporary location for **Enter the folder name on the selected datastore which has pre-existing virtual disks**.

In order for Double-Take Availability to find the pre-existing disk, the virtual disk file names must be formatted using the convention SourceServer_ DriveLetter. For example, if your source server is Alpha and you are protecting drives C and D, Double-Take Availability will look for the file names Alpha_C.vmdk and Alpha_D.vmdk. If you are using IP addresses, substitute the IP address for the server name. For example, if the IP address for server Alpha is 172.31.10.25 then Double-Take Availability will look for the file names 172.31.10.25_C.vmdk and 172.31.10.25_D.vmdk.

If you originally created a virtual disk and specified the source server by its IP address, the pre-existing virtual disk file name cannot use the server name. However, you can rename that file and its associated -flat.vmdk file to use the IP address. The reverse is also true. If you originally specified the source server by its name, the pre-existing virtual disk file name cannot use the server's IP address. However, you can rename the file and its associated -flat.vmdk to use the source name. For example, if you originally created a virtual disk and specified the source by its IP address, you need to rename the file source_name_drive.vmdk to source_IPaddress_drive.vmdk. You also need to rename the file source_name_drive-flat.vmdk to source_IPaddress_ drive-flat.vmdk. The reverse (change source_IPaddress to source_name for both files) is also true. Additionally, you will need to edit the .vmdk file manually because it contains the name of the -flat.vmdk file. Modify the reference to the -flat.vmdk file to the new name you have specified using any standard text editor.

In a WAN environment, you may want to take advantage of the **Use pre-existing virtual disks** feature by using a process similar to the following.

a. Create a protection job in a LAN environment, letting Double-Take Availability create the virtual disk for you.
b. Complete the mirror process locally.
c. Disconnect the protection job.
d. Shut down and move the ESX target server to your remote site.

e. After the target server is back online at the remote site, create a new protection job for the same source server and select to **Use pre-existing virtual disks**. Double-Take Availability will reuse the existing .vhd files and perform a difference mirror over the WAN to bring the virtual machine up-to-date.

15. Click **Next** to continue.
16. Configure the replica virtual machine.

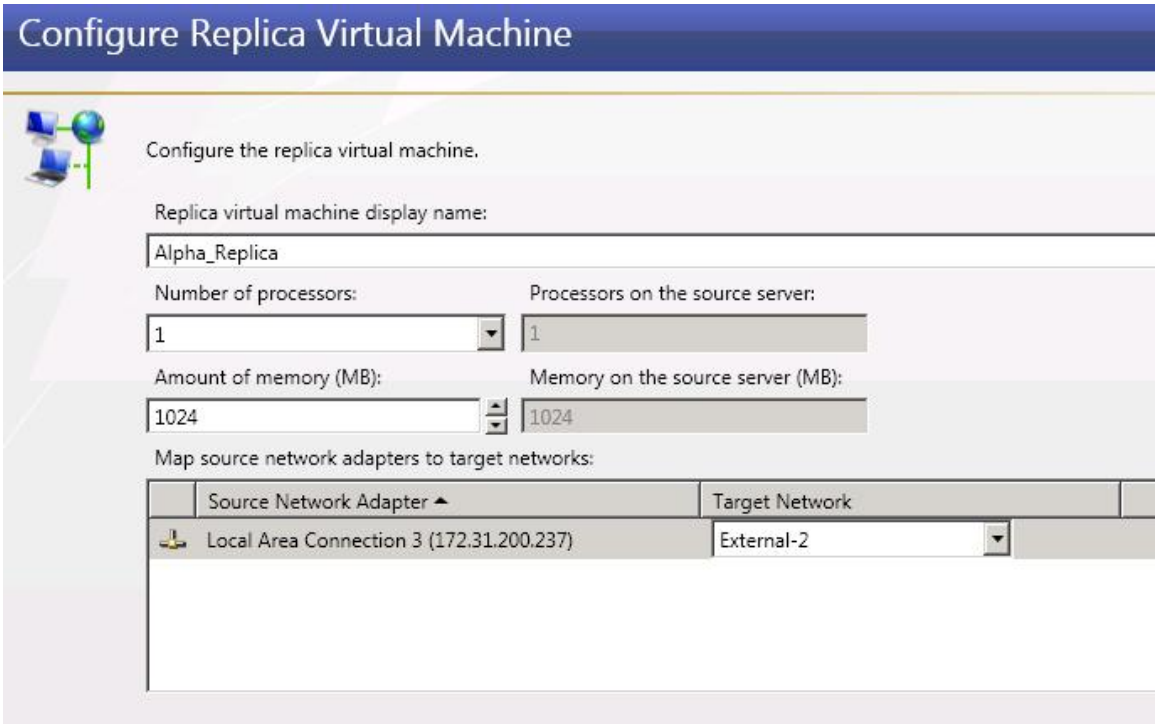

- **Replica virtual machine display name**—Specify the name of the replica virtual machine. This will be the display name of the virtual machine on the host system.
- **Number of processors**—Specify how many processors to create on the new virtual machine. The number of processors on the source is displayed to guide you in making an appropriate selection. If you select fewer processors than the source, your clients may be impacted by slower responses.
- **Amount of memory**—Specify the amount of memory, in MB, to create on the new virtual machine. The memory on the source is displayed to guide you in making an appropriate selection. If you select less memory than the source, your clients may be impacted by slower responses.
- **Map source network adapters to target network adapters**—Identify how you want to handle the network mapping after failover. The **Source Network Adapter** column lists the NICs from the source. Map each one to a **Target Network**, which is a virtual network on the target.

17. Click **Next** to continue.
18. Specify your protection settings.

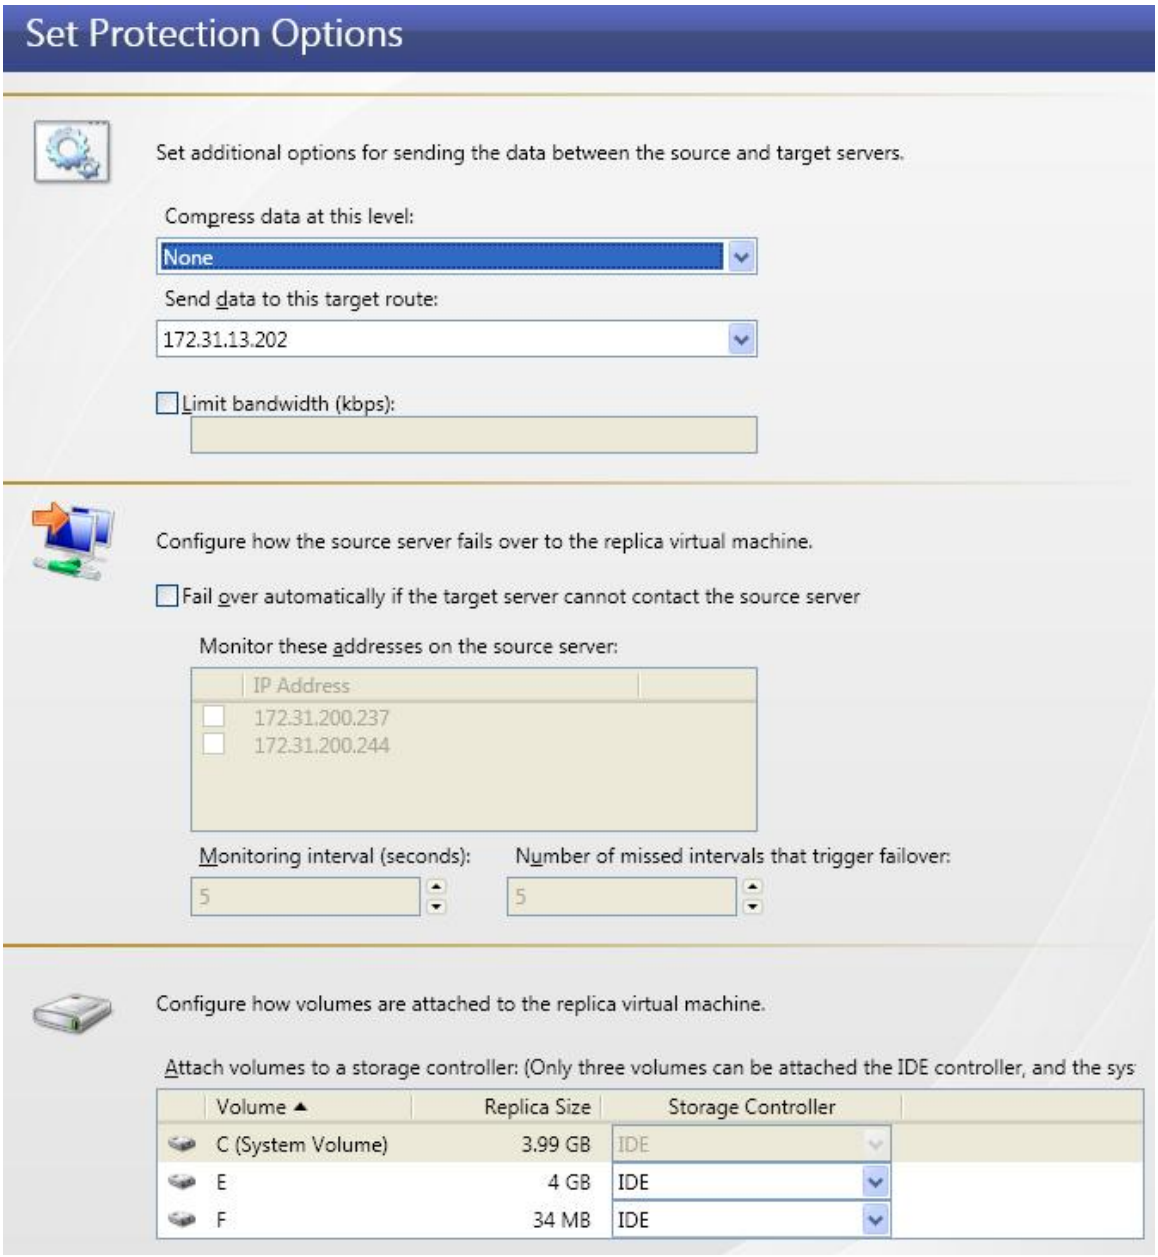

The fields on this screen will vary depending on if you are using an ESX target or a Hyper-V target.

- **Compress data at this level**—Specify the level of compression that you want to use for your transmissions from the source to the target. If compression is enabled, the data is compressed before it is transmitted from the source. When the target receives the compressed data, it decompresses it and then writes it to disk.

- **Send data to this target route**—By default, Double-Take Availability will select a default target route for transmissions. If desired, select a different target route for transmissions.
- **Limit bandwidth**—Bandwidth limitations are available to restrict the amount of network bandwidth used for Double-Take Availability data transmissions. When a bandwidth limit is specified, Double-Take Availability never exceeds that allotted amount. The bandwidth not in use by Double-Take Availability is available for all other network traffic. If desired, enter a value, in kilobits per second, to limit data transmission. The value you enter is the maximum amount of data that will be transmitted per second.
- **Fail over automatically if the target server cannot contact the source server**—If this option is selected, failover will automatically occur when the target can no longer contact the source. If this option is disabled, you will have to monitor communications between the target and source and manually initiate failover when the target can no longer contact the source. In either case, failover is not available until after the initial mirror has been completed.
- **Monitor these addresses on the source server**—If you have enabled automatic failover, specify the IP addresses on the source that the target should monitor. Failover will be triggered when one of the monitored IP addresses is identified as failed.
- **Monitoring interval**—Specify the number of seconds between monitor requests sent from the target to the source to determine if the source is online.
- **Number of missed intervals that trigger failover**—Specify the number of monitor replies sent from the source to the target that can be missed before assuming the source has failed.

Note: To achieve shorter delays before failover, use lower interval and missed interval values. This may be necessary for servers, such as a web server or order processing database, which must remain available and responsive at all times. Lower values should be used where redundant interfaces and high-speed, reliable network links are available to prevent the false detection of failure. If the hardware does not support reliable communications, lower values can lead to premature failover. To achieve longer delays before failover, choose higher values. This may be necessary for servers on slower networks or on a server that is not transaction critical. For example, failover would not be necessary in the case of a server restart.

- **Attach volumes to a storage controller**—If you are using a Hyper-V target, you can specify the type of **Storage Controller** that you want to use for each volume on the Hyper-V target.

Note: A maximum of three volumes can be attached to an IDE controller. If you are protecting more than three volumes on the source, you will need to install the Hyper-V Integration Components to acquire a SCSI device. See your Hyper-V documentation for more information.

If your source is Windows 2003 or Windows 2008 with no service packs and you have selected a SCSI controller, you will need to manually install the Hyper-V Integration Components after failover to attach these volumes to the replica virtual machine.

---

19. Click **Next** to continue.
20. The **Protection Summary** page displays all of the options you selected in your workflow. If you want to make any changes to any of the workflow settings, click **Back** to return to previous pages of the workflow. If you want to modify the name assigned to this protection job, click **Change** and specify a new name.
21. When you are satisfied with your workflow selections, click **Finish**, and you will automatically be taken to the Monitor Connections page.

Protecting a Hyper-V server to a Hyper-V server

Use these instructions to protect a virtual machine on a Hyper-V server, where you want to protect the host-level virtual disk files (the .vhd files), to a virtual server on a Hyper-V server.

1. Open the Double-Take Console by selecting **Start**, **Programs**, **Double-Take**, **Double-Take Console**.
2. Click **Get Started** from the toolbar.
3. Select **Double-Take Availability** and click **Next**.
4. Select **Protect a Hyper-V virtual machine using host-level protection** and click **Next**.
5. Specify your source server. This is the Hyper-V source that contains the virtual machine that you want to protect.

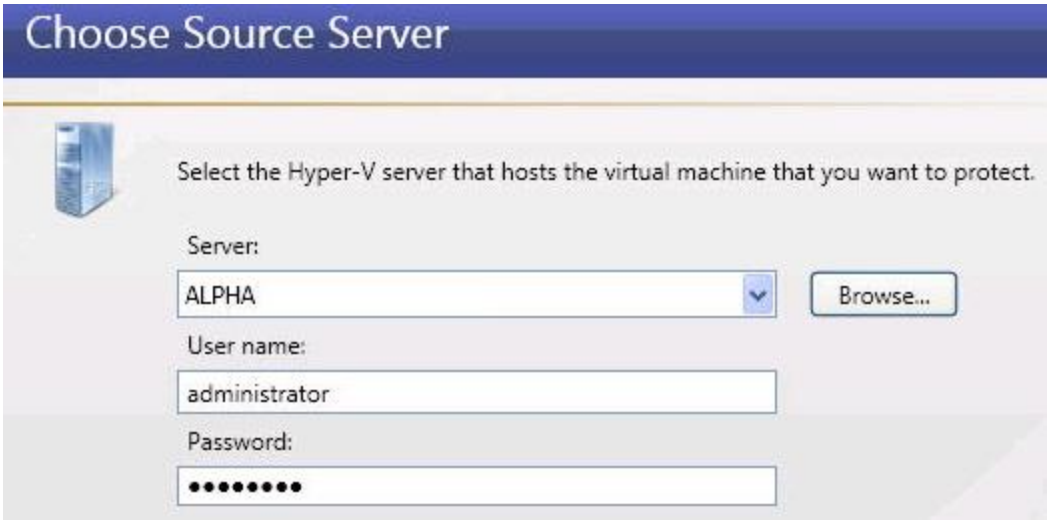

- **Server**—Specify the name or IP address of the Hyper-V server that is hosting the virtual machine that you want to protect. You can click **Browse** to select a server from a network drill-down list.

Note: If you enter the source server's fully-qualified domain name, the Double-Take Console will resolve the entry to the server short name. If that short name resides in two different domains, this could result in name resolution issues. In this case, enter the IP address of the server.

- **User name**—Specify a user that is a member of the Double-Take Admin and local administrator security groups on the source. If you want to use a domain user account, enter a fully-qualified domain name in the format domain\username or username@domain.
- **Password**—Specify the password associated with the **User name** you entered.

6. Click **Next** to continue.
7. Choose the virtual machine on the Hyper-V source that you want to protect. Select a virtual machine from the list and click **Next** to continue.

8. Specify the Hyper-V target server where you will store the replica of the source server.

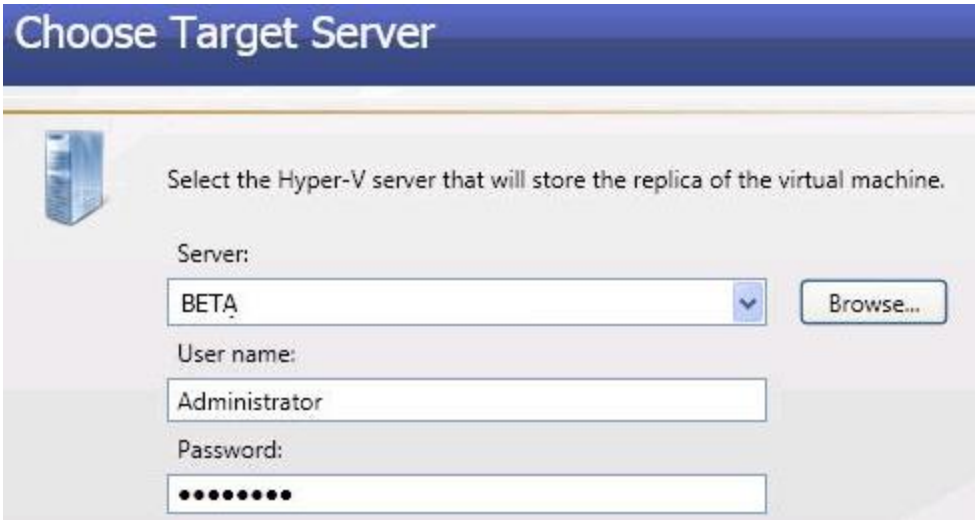

- **Server**—Specify the name or IP address of the Hyper-V target server. You can click **Browse** to select a server from a network drill-down list.

Note: If you enter the target server's fully-qualified domain name, the Double-Take Console will resolve the entry to the server short name. If that short name resides in two different domains, this could result in name resolution issues. In this case, enter the IP address of the server.

- **User name**—Specify a user that has administrative rights for Microsoft Hyper-V and is a member of the Double-Take Admin security group. If you want to use a domain user account, enter a fully-qualified domain name in the format domain\username or username@domain.
- **Password**—Specify the password associated with the **User name** you entered.

9. Click **Next** to continue.
10. Select a home folder on the Hyper-V target for the replica virtual machine.

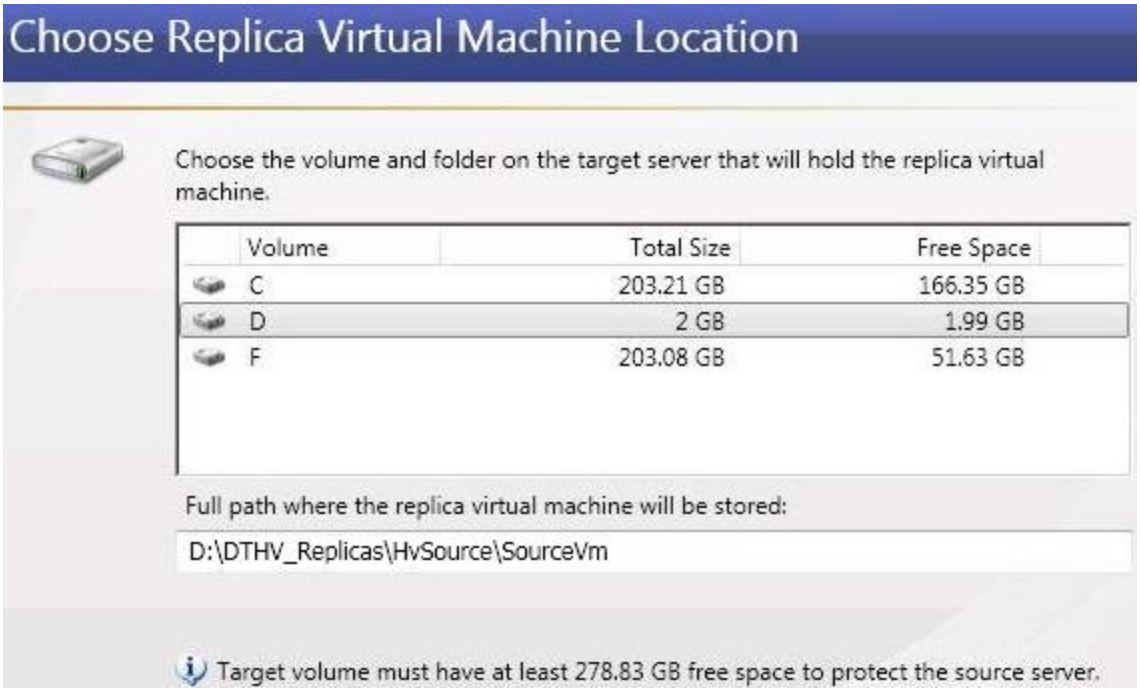

- **Volume**—Select one of the volumes from the list to indicate which volume on the target where you want to store the new virtual server when it is created. The target volume must have enough **Free Space** to store the source data. The minimum size is noted at the bottom of the page.
- **Full path where the replica virtual machine will be stored**—Specify a location on the selected **Volume** to store the replica of the source. By

specifying an existing folder, you can reuse an existing virtual machine on your Hyper-V target created by a previous protection job. This can be useful for pre-staging data on a virtual machine over a LAN connection and then relocating it to a remote site after the initial mirror is complete. When you reuse a virtual machine, Double-Take Availability performs a difference mirror which saves time. Use the following steps to reuse a virtual machine on a Hyper-V target.

a. Create a protection job in a LAN environment, letting Double-Take Availability create the virtual disk for you.
b. Complete the mirror process locally.
c. Disconnect and delete the protection job and when prompted, select to **Keep and register the replica virtual machine**.
d. From the Hyper-V manager, delete the replica virtual machine, which will delete the virtual machine configuration but will keep the associated .vhd files.
e. Shut down and move the Hyper-V target server to your remote site.
f. After the target server is back online at the remote site, create a new protection job for the same source server and select to reuse the existing virtual machine when prompted. Double-Take Availability will reuse the existing .vhd files and perform a difference mirror over the WAN to bring the virtual machine up-to-date.

11. Click **Next** to continue.
12. Configure the replica virtual machine.

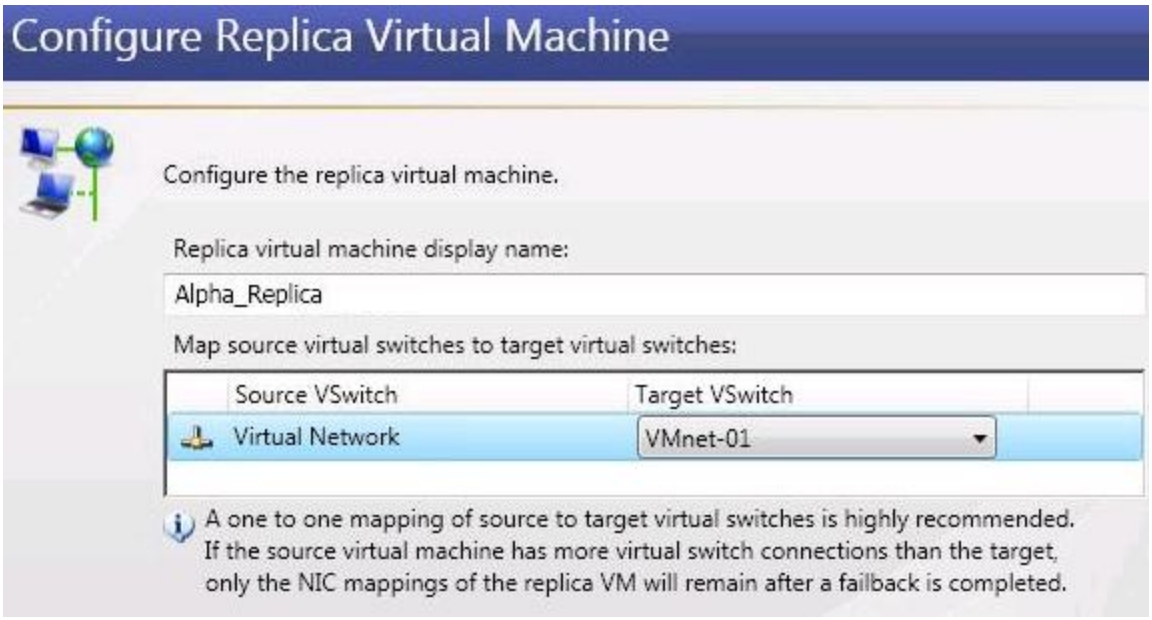

- **Replica virtual machine display name**—Specify the name of the replica virtual machine. This will be the display name of the virtual machine on the host system.
- **Map source virtual switches to target virtual switches**—Identify how you want to handle the network mapping after failover. The **Source VSwitch** column lists the virtual networks from the source. Map each one to a **Target VSwitch**, which is a virtual network on the target.

13. Click **Next** to continue.
14. Specify your protection settings.

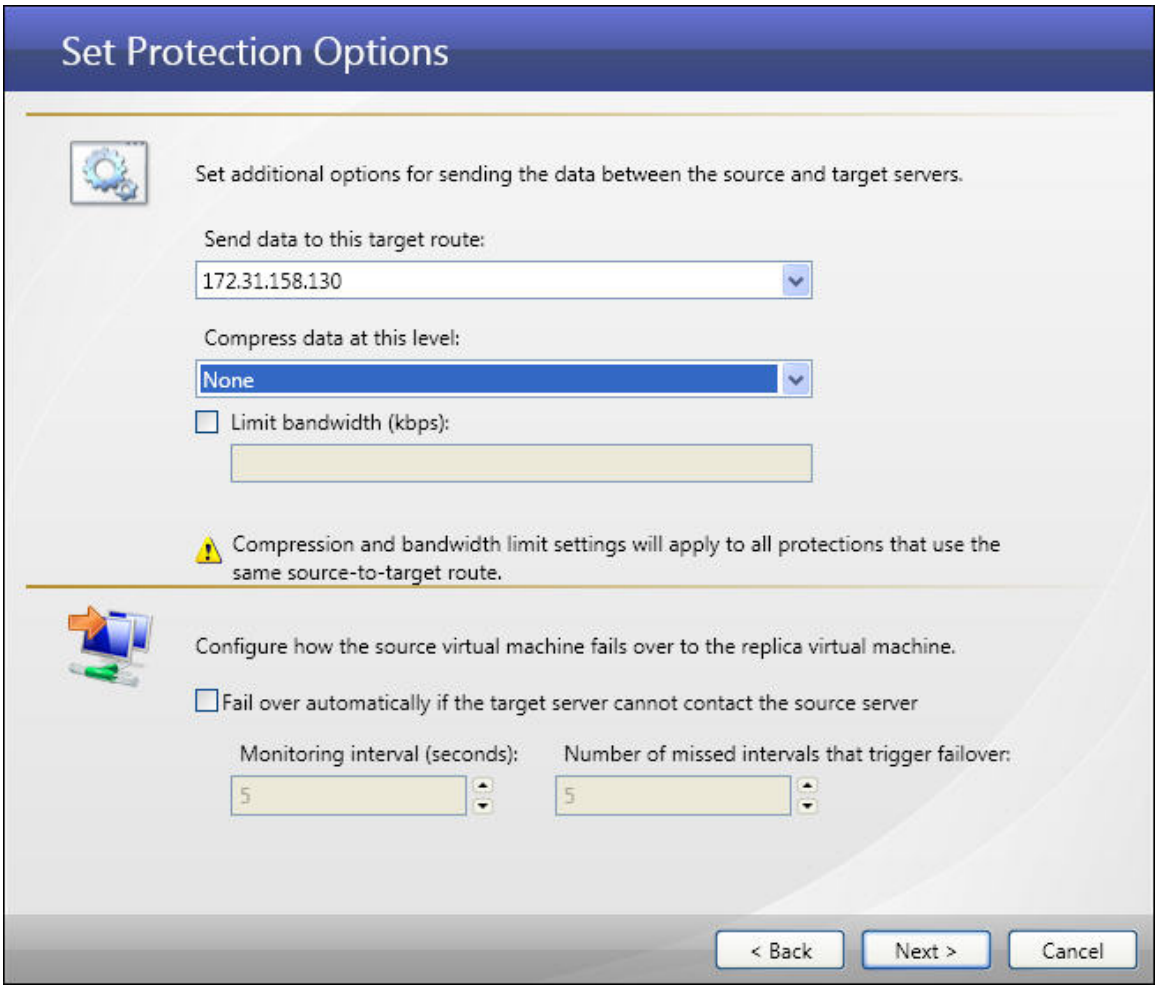

- **Send data to this target route**—By default, Double-Take Availability will select a default target route for transmissions. If desired, select a different target route for transmissions.
- **Compress data at this level**—Specify the level of compression that you want to use for your transmissions from the source to the target. If compression is enabled, the data is compressed before it is transmitted from

the source. When the target receives the compressed data, it decompresses it and then writes it to disk.

- **Limit bandwidth**—Bandwidth limitations are available to restrict the amount of network bandwidth used for Double-Take Availability data transmissions. When a bandwidth limit is specified, Double-Take Availability never exceeds that allotted amount. The bandwidth not in use by Double-Take Availability is available for all other network traffic. If desired, enter a value, in kilobits per second, to limit data transmission. The value you enter is the maximum amount of data that will be transmitted per second.
- **Fail over automatically if the target server cannot contact the source server**—If this option is selected, failover will automatically occur when the target can no longer contact the source. If this option is disabled, you will have to monitor communications between the target and source and manually initiate failover when the target can no longer contact the source. In either case, failover is not available until after the initial mirror has been completed.
- **Monitoring interval**—Specify the number of seconds between monitor requests sent from the target to the source to determine if the source is online.
- **Number of missed intervals that trigger failover**—Specify the number of monitor replies sent from the source to the target that can be missed before assuming the source has failed.

Note: To achieve shorter delays before failover, use lower interval and missed interval values. This may be necessary for servers, such as a web server or order processing database, which must remain available and responsive at all times. Lower values should be used where redundant interfaces and high-speed, reliable network links are available to prevent the false detection of failure. If the hardware does not support reliable communications, lower values can lead to premature failover. To achieve longer delays before failover, choose higher values. This may be necessary for servers on slower networks or on a server that is not transaction critical. For example, failover would not be necessary in the case of a server restart.

---

15. Click **Next** to continue.
16. The **Protection Summary** page displays all of the options you selected in your workflow. If you want to make any changes to any of the workflow settings, click

Back to return to previous pages of the workflow. If you want to modify the name assigned to this protection job, click **Change** and specify a new name.

17. When you are satisfied with your workflow selections, click **Finish**, and you will automatically be taken to the Monitor Connections page.

Note: Once protection is established, Double-Take Availability monitors the virtual disks of the protected virtual machine for changes to the disk layout. If a new virtual hard disk is added to the virtual machine, the protection job will automatically be updated to include the new virtual hard disk, and a file difference mirror will automatically start. However, if a virtual hard disk is removed from the protected virtual machine, the virtual hard disk will not be removed from the projection job until it is deleted from the source or the protection job is deleted and re-created.

Protecting an ESX server to an ESX server

Use this process to protect a virtual machine on an ESX server, where you want to protect the host-level virtual disk files (the .vmdk files), to a virtual server on an ESX server. In this scenario, you have several steps to complete.

1. Port configuration—You must configure your ESX servers, your VirtualCenter server(s), and the machine running the Double-Take Availability for VMware Infrastructure service. These steps only need to be performed once.
2. Root or non-root login configuration—You must configure your ESX servers for either root or non-root login. These steps only need to be performed once.
3. Establish protection—Use these instructions to configure your ESX to ESX protection job.
4. Optional ESX protection settings—Optional settings can be configured at the end of your protection job creation or after the job has been established.

Configuring ports

1. Because the Double-Take Availability for VMware Infrastructureconsole uses SSH to communicate with the VMware ESX host(s), and the ESX hosts also use it to connect to each other, you must configure SSH port communication on both your source and target ESX servers.
    a. Using the VMware Virtual Infrastructure Client, select the host ESX server.
    b. On the **Configuration** tab, select **Security Profile**.
    c. In the **Firewall Properties**, verify that **SSH Client** is selected and click **OK**.
    d. Verify that the **Configuration** tab shows that **SSH Client**is enabled for the host on port 22.
    e. In the **Firewall Properties**, verify that **SSH Server** is selected and click **OK**.
    f. Verify that the **Configuration** tab shows that **SSH Cerver**is enabled for the host on port 22.
2. Open port 443 on both your source and target ESX servers for HTTPS communication. You also need to open this port on your VirtualCenter server.
3. Open port 6331 on the machine running the Double-Take Availability for VMware Infrastructure service.

Note: If the Double-Take Availability for VMware Infrastructureconsole and the Double-Take Availability for VMware Infrastructure service are on separate machines, the machines cannot be separate by a firewall because Microsoft .NET remoting is required, which is not compatible with a firewall configuration.

Configuring root or non-root login

You must configure both your source and target ESX servers to allow either root or non-root login. If you are not using VirtualCenter, you must use root account credentials. If you want to use non-root credentials, VirtualCenter is required. In addition to using non-root credentials, VirtualCenter allows you to use VMotion to move the virtual machine.

- **Root account login**—If you are not using VirtualCenter, you must use root account credentials.
    1. Login to the host ESX server using root credentials.
    2. Using a text editor, open the file /etc/ssh/sshd_config.
    3. Locate the line PermitRootLogin no and change it to PermitRootLogin yes.
    4. Save the configuration file.
    5. From a command line, enter the following command service sshd restart, which will restart the SSH service.
- **Non-root account login**—If you want to use non-root credentials, VirtualCenter is required. In addition to using non-root credentials, VirtualCenter allows you to use VMotion to move the virtual machine.
    1. Login to the host ESX server as root.
    2. Use the sudo command to transition to the root.
    3. Use the adduser <username> command to create a new user.
    4. Execute visudo to modify the sudo configuration file (/etc/sudoers). Do not modify the file directly.
    5. Add the following line to the configuration file.

       username ALL=(ALL) NOPASSWD: ALL
    6. Save the file and exit (:w!).
    7. Logout and log back in as a sudo account.
    8. Execute the following command to make sure you can access VMware datastores on /vmfs/volumes. If the command succeeds (prints the counts of the root home folder), then sudo is configured correctly. If it fails with a permission denied error or prompts for a password, then sudo is misconfigured.

       sudo ls ~root

Establishing ESX to ESX protection

1. Open the Double-Take Availability for VMware Infrastructure console by selecting **Start**, **Programs**, **Double-Take**, **Availability**, **Double-Take for VMware Infrastructure**.
2. If prompted, connect to the Double-Take Availability for VMware Infrastructure server. This prompt will appear if this is the first time you have accessed the console, if there are no valid, saved credentials, or if you have intentionally disconnected from a server.
    - **Server**—Specify the name or IP address of the Double-Take Availability for VMware Infrastructure server.
    - **User name**—Specify a user in the local administrator group that will access the Double-Take Availability for VMware Infrastructure server.
    - **Password**—Specify the password associated with the **User name** you entered.
3. If you want to save the login information so that you do not have to login to the server the next time you use the console, enable **Save DTAVI connection information**.
4. Click **Connect** to connect to the server.
5. To begin the protection workflow, select **Protect a virtual machine**.
6. Select the virtual machine that you want to protect.

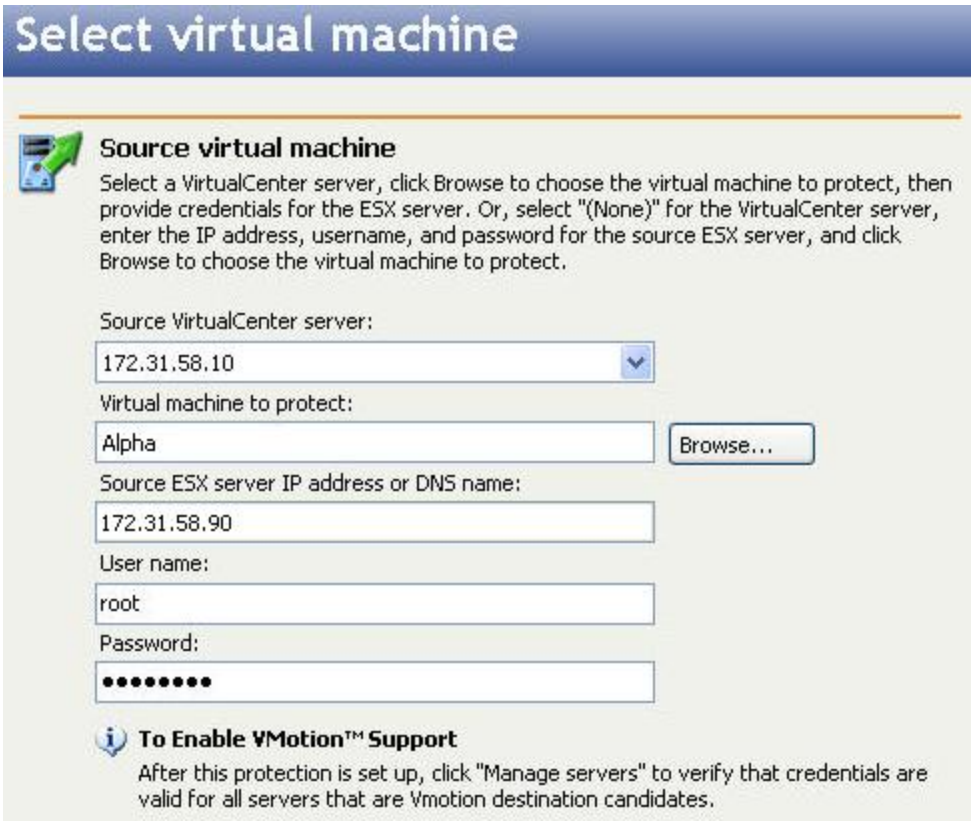

- **Source VirtualCenter server**—Select the VirtualCenter server that administrators the virtual machine you want to protect. If you are not using VirtualCenter, select **None**.
- **Virtual machine to protect**—Specify the name of the virtual server you want to protect. You can click **Browse** to select a server from a network drill-down list or to search for a server. Once you have specified a virtual server to protect, the IP address or DNS of the ESX server will be displayed.

Note: Each protection job applies to a single virtual machine.

Virtual machines using ESX raw or independent disks are not supported.

If your virtual machine is configured to use thin (sparse disks), the replica on the target will not be a thin disk. Make sure there is adequate space on the target.

The virtual machine name cannot contain any of the following special characters.

/ \ : * ? ' " < > |

- **User name**—Specify the user account that will log in to the ESX server.
- **Password**—Specify the password associated with the **User name** you entered.

7. Click **Next** to continue.
8. Specify the target ESX server that will store the replica virtual machine.

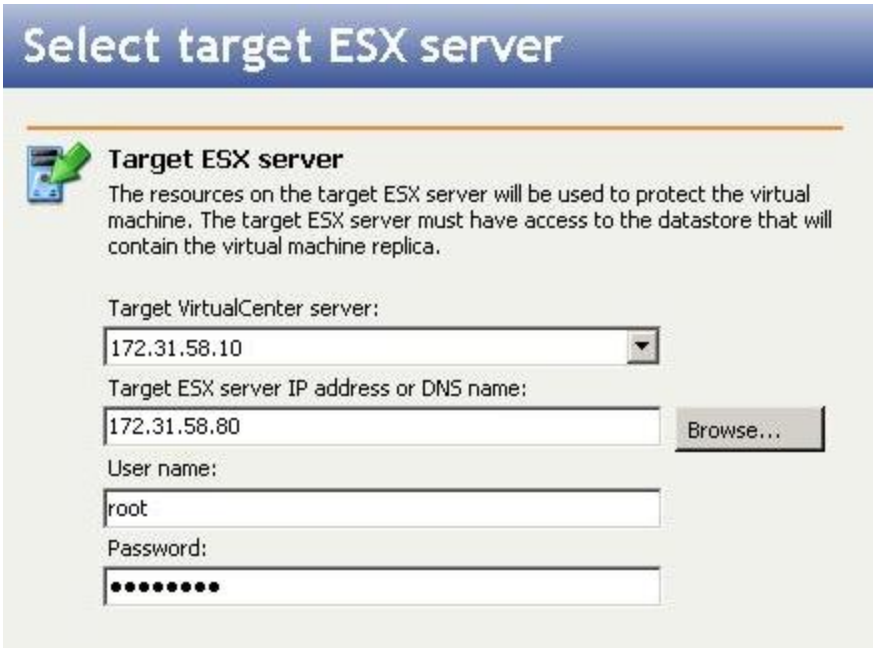

- **Target VirtualCenter server**—Select the VirtualCenter server that administrators the target replica virtual machine. If you are not using VirtualCenter, select **None**
- **Target ESX server IP address or DNS name**—Enter the IP address or DNS name of the ESX server that will host the target replica virtual machine. You can click **Browse** to select a server from a network drill-down list or to search for a server.
- **User name**—Specify the user account that will log in to the ESX server.
- **Password**—Specify the password associated with the **User name** you entered.

9. Click **Next** to continue.
10. Select a datastore on the target where the source virtual machine data will be replicated.

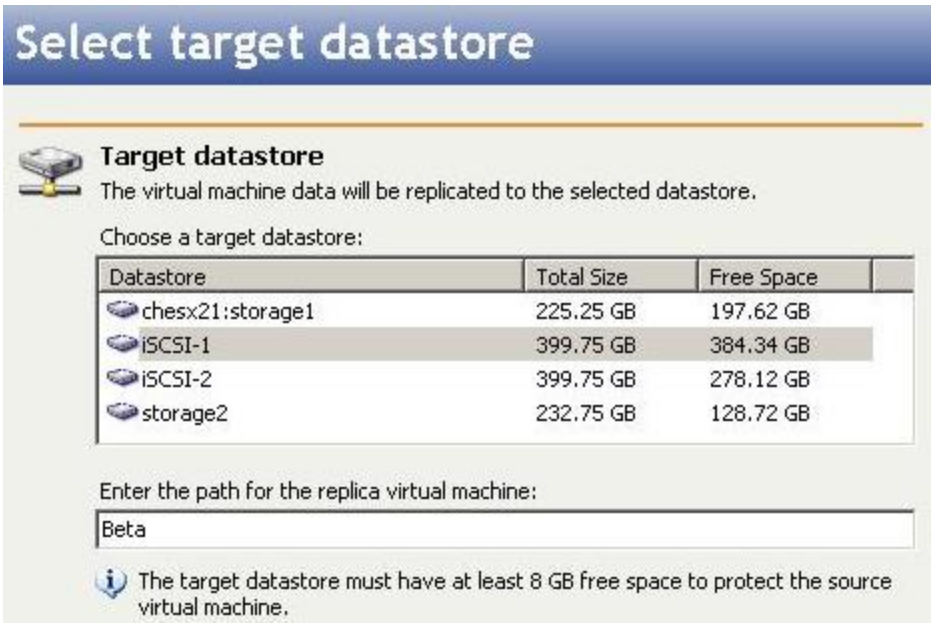

- **Choose a target datastore**—Select a datastore in the table that is large enough to hold the source virtual machine. The minimum size is noted at the bottom of the page.
- **Enter the path for the replica virtual machine**—Specify a location to store the replica of the source virtual machine. By default, the location will be named the source virtual machine name.

Note: The replica virtual path cannot contain any of the following special characters.

/ \ : * ? ' " < > |

By specifying a datastore and path with an existing virtual disk, you can reuse an existing virtual machine created by a previous protection job. This can be useful for pre-staging data on a virtual machine over a LAN connection and then relocating it to a remote site after the initial mirror is complete. When you reuse a virtual machine, Double-Take Availability for VMware Infrastructure performs a difference mirror which saves times. Use the following steps to reuse a virtual machine.

1. Verify the source has at least one active snapshot, thus unlocking the .vmdk files allowing them to be copied.
2. Create a protection job, but delay the protection start time. Provide a long enough delay to copy the .vmdk files from the source to the target.

3. Click **View replica virtual machine disk mapping** to determine the proper location on the target to copy the .vmdk files.
4. Copy the source .vmdk files to the target using the exact locations and filenames specified in the mapping file. Make sure that your copy method does not modify the size of the file. The size of the .vmdk files on the source and target must match.
5. Remove the original snapshot.
6. Modify the protection start time to begin immediately. Double-Take Availability for VMware Infrastructure will reuse the .vmdk files that you copied to the target and perform a difference mirror to bring the files up-to-date.

11. Click **Next** to continue.
12. Configure the replica virtual machine.

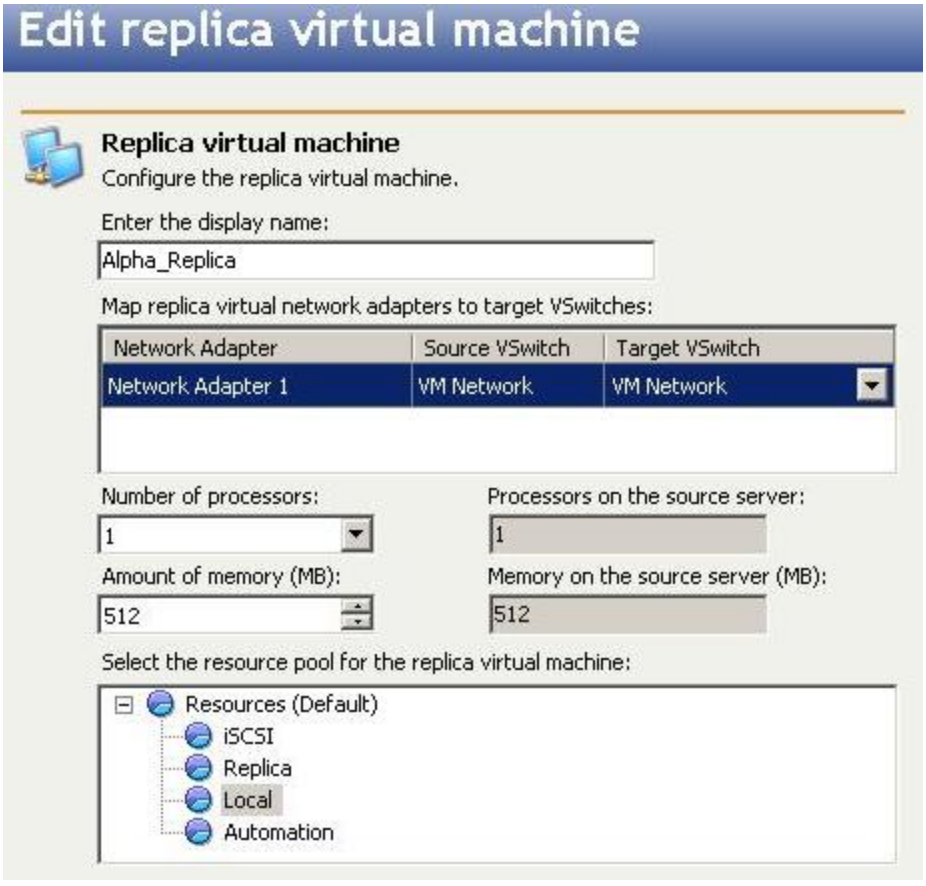

- **Enter the display name**—Specify the name for the replica virtual machine. The name cannot contain any of the following special characters.

  # / \ : * ? " " < > |

- **Map replica virtual network adapters to target VSwitches**—Identify how you want to handle the network mapping after failover. The **Source VSwitch** column lists the virtual networks from the source. Map each one to a **Target VSwitch**, which is a virtual network on the target.
- **Number of processors**—Specify how many processors you want on the replica virtual machine. The number of processors on the source is displayed to guide you in making an appropriate selection. If you select fewer processors than the source, your clients may be impacted by slower responses.
- **Amount of memory**—Specify the amount of memory, in MB, you want on the replica virtual machine. The memory on the source is displayed to guide you in making an appropriate selection. If you select less memory than the source, your clients may be impacted by slower responses.
- **Select the resource pool for the replica virtual machine**—Specify the resource pool to allocate host-provided CPU and memory to the replica virtual machine.

13. Click **Next** to continue.
14. The **Protection Summary** page displays the options you selected in your workflow, as well as optional protection settings.
    - If you want to make any changes to any of the workflow settings, click **Back** to return to previous pages of the workflow.
    - If you want to delay when the protection job is enabled, click **Schedule**.
    - If you want to modify the name assigned to this protection job, click **Change** and specify a new name.
    - If you want to modify any of the optional settings, click **Change** next to the setting.
15. When you are satisfied with your workflow selections, click **Finish**, and you will automatically be taken to the Monitor protection page.

Note: The target virtual machine is registered when replication is started and will remain registered. To unregister a machine, you must click **Delete Protection** and choose **Delete the associated replica virtual machine**.

Even though the target virtual machine appears to be available on the ESX server, it should not be powered on, removed, or modified while it is owned by an active protection job, otherwise the target virtual machine will become corrupt and break the protection job.

Do not attempt to manually create or delete snapshots on the protected virtual machine. This will disrupt the protection of the virtual machine and

may generate unpredictable results on the source and target virtual machines.

Double-Take Availability for VMware Infrastructure supports generic SCSI device mappings in virtual machines, however, the generic SCSCI device will not be created on the target virtual machine during failover because the target virtual machine will fail to start if the SCSI device hardware does not exist on the target ESX host. After failback, the generic SCSI device will be mapped back to the original source, but the SCSI device will not be mapped back to the original source if the protection job is re-created in the reverse direction.

Optional ESX protection settings

Optional protection settings are available when configuring an ESX to ESX protection job, but they are not required. The options are available at the end of the workflow when you are establishing protection or from the **Monitor protection** page when you click **Configure Protection**.

- Scheduling protection
- Changing the name of the protection job
- Setting transmission options
- E-mailing notifications
- Updating VirtualCenter credentials
- Configuring restart and threshold options

Scheduling protection

1. From the **Protection summary** page, click **Schedule**.
2. If you want to start the protection immediately, select **Start this protection immediately**.
3. If you want to delay the start of the protection, select **Schedule this protection to start**.
4. Select a date and time for the protection job to start.
5. Click **Save**.

Changing the name of the protection job

1. From the **Protection summary** page, click **Change** in the **Name** section.
2. Specify a new name for the protection job.
3. Click **Save**.

Setting transmission options

1. From the **Protection summary** page, click **Change** in the **Data transmission** section.
2. Specify any of the following data transmission options.
   - **Chose when to use compression**—Specify the level of compression you want to use. Compression reduces the amount of bandwidth needed to transmit data from the source to the target. The data is compressed before being transmitted and then is uncompressed before it is written on the target. Typically, compression is used in WAN environments, but not in LAN environments. If desired, enable compression and select the level of compression that you want to use. All connections to the same target will have the same compression settings.
   - **Transmit when the snapshot data reaches this size**—Specify the size, in MB, that will trigger when snapshots of the source are transmitted to the target. When multiple virtual disks are used, any combination of writes across all virtual disks that accumulate to the specified size will trigger transmission.
   - **Transmit data, regardless of snapshot size, after**—Specify a length of time that will trigger when snapshots of the source are transmitted to the target. Transmission will occur regardless of the size of the snapshots.

     Note: A snapshot transmission cycle will begin when either of the time or size threshold conditions are met. During the snapshot transmission cycle, the thresholds are not monitored. After the snapshot transmission cycle has completed, the application will again monitor the thresholds. If either of the thresholds were crossed during the snapshot transmission cycle, a new transmission cycle will begin immediately.

     You may want to adjust the snapshot transmission options to optimize performance in your environment. Some factors you need to consider when adjusting these settings include the volume of write traffic in the virtual machine, the allowed data loss time period, and the cost to the virtual infrastructure.

   - **Send data to this route**—By default, Double-Take Availability for VMware Infrastructure will select the default route for transmissions between the two ESX servers. If desired, select a different IP address on the target that will be used for transmissions.

- **Limit bandwidth**—Bandwidth limitations are available to restrict the amount of network bandwdith used for Double-Take Availability for VMware Infrastructure data transmissions. When a bandwidth limit is specified, Double-Take Availability for VMware Infrastructure never exceeds that allotted amount. The bandwidth not in use by Double-Take Availability for VMware Infrastructure is available for all other network traffic. Enter a value in kilobits per second to limit data transmission. This is the maximum amount of data that will be transmitted per second.

3. Click **Save**.

E-mailing notifications

1. From the **Protection summary** page, click **Change** in the **E-mail notifications** section.

Note: In order to use e-mail notification, you may need to complete any or all of the following on the Double-Take Availability for VMware Infrastructure server.

- Disable anti-virus software
- Open port 25 in your anti-virus software to allow SMTP e-mail
- Enable outbound e-mail messages
- Exclude VI_Service.exe from blocked processes for sending outbound e-mail messages.

---

2. If you have not yet configured an e-mail server, you will be prompted at the top of the window. Click **Configure**.
3. Specify the e-mail notification settings.
   - **Recipients**—Enter the e-mail addresses where the e-mail messages should be sent. Separate the addresses with a comma, semi-colon, or carriage return.
   - **Notifications**—Select the event categories that you want to be notified about. If there are no event categories selected, there will be no e-mail notifications.
4. Click **Test E-mail Settings** to verify your e-mail configuration.

Note: The following error message indicates anti-virus software on the server may be blocking outbound e-mail messages.

> Failure sending mail. Unable to connect to the remote server. An established connection was aborted by the software in your host machine.

---

5. Click **Save**.

Updating VirtualCenter credentials

1. If you configured VirtualCenter servers when you established protection, you can select **Change** in the **VirtualCenters** section on the **Protection Summary** page.
2. Modify the User name and Password associated with either the source or target VirtualCenter server. The changes will be applied to this protection job only.
3. Click **Save**.

Configuring restart and threshold options

1. From the **Protection summary** page, click **Change** in the **Restart and thresholds** section.
2. Specify any of the following options.
   - **Restart this protection automatically if there is a problem**—This option specifies if Double-Take Availability for VMware Infrastructure will attempt to restart the protection job if there is a problem. If you enable this option, specify the **Number of times to attempt to restart**.
   - **Disk space remaining**—Specify the amount of space remaining, either by percentage or by size in MB, to trigger stopping the protection.
3. Click **Save**.

Using firewalls with virtual workloads

If your source and target are on opposite sides of a firewall, you will need to configure your hardware to accommodate communications. You must have the hardware already in place and know how to configure the hardware ports. If you do not, see the reference manual for your hardware.

These firewall instructions are for the following protection scenarios.

- Your source is a physical or virtual server, you want to protect the volumes from the physical server or the volumes from within the virtual guest operating system, and your target is a virtual server on a Hyper-V server.
- Your source is a physical or virtual server, you want to protect the volumes from the physical server or the volumes from within the virtual guest operating system, and your target is a virtual server on an ESX server.
- Your source is a Hyper-V virtual server, you want to protect the host-level virtual disk files (the .vhd files), and your target is a virtual server on a Hyper-V server.

- Virtual workload ports
- Microsoft Windows ports
- Hardware ports

Virtual workload ports

By default, Double-Take Availability uses port 6320 for all communications. To verify or modify the ports, you must use the Replication Console.

1. Open the Replication Console.
2. Locate your servers in the server tree in the left pane of the Replication Console.
3. Right-click the server in the left pane of the Replication Console and select **Properties**.
4. On the **Network** tab, verify or modify the **Communications Port** as needed. All servers and clients in your virtual workload protection must be using the same port.

Double-Take Availability uses ICMP pings to monitor the source for failover. A failover monitor will not be created if ICMP is blocked (although the data and system state will still be protected). You should configure your hardware to allow ICMP pings between the source and target. If you cannot, you will have to monitor for a failure manually and create a dummy monitor at failover time that can be manually failed over. Contact technical support for assistance with this manual process.

Microsoft Windows ports

Virtual workload protection uses WMI (Windows Management Instrumentation) which uses RPC (Remote Procedure Call). By default, RPC will use ports at random above 1024, and these ports must be open on your firewall. RPC ports can be configured to a specific range by specific registry changes and a reboot. See the Microsoft Knowledge Base article 154596 for instructions.

Virtual workload protections also rely on other Microsoft Windows ports.

- Microsoft File Share uses ports 135 through 139 for TCP and UDP communications.
- Microsoft Directory uses port 445 for TCP and UDP communications.

These ports must be open on your firewall. Check your Microsoft documentation if you need to modify these ports.

Hardware ports

You need to configure your hardware so that the Double-Take Availability ports and Microsoft Windows ports are open. Since communication occurs bidirectionally, make sure you configure both incoming and outgoing traffic.

There are many types of hardware on the market, and each can be configured differently. See your hardware reference manual for instructions on setting up your particular router.

Cluster protection

Your cluster protection will depend on the type of cluster configuration you are using.

- **Standard cluster configuration**—If you are using a standard cluster, where a single copy of data resides on a SCSI disk that is shared between cluster nodes, you will be using the Double-Take Source Connection cluster resource to protect your cluster.
- **GeoCluster configuration**—If you are using a GeoCluster configuration, where data is stored on volumes local to each node and replicated to each node in the cluster, you will be using the GeoCluster Replicated Disk cluster resource to protect your cluster.

Protecting a standard cluster

Use the following steps to protect a standard cluster, where a single copy of data resides on a SCSI disk that is shared between cluster nodes. You will be using the Double-Take Source Connection cluster resource to protect your standard cluster.

1. If your source is a cluster, create a virtual server (including resources for an IP address, network name, and physical disk) on the source cluster. With this configuration, users will access their data from the source cluster, regardless of which node is currently in control. MSCS will handle failover between the nodes of the cluster. Double-Take Availability will handle failover between the source cluster and the target (cluster or standalone). See your Microsoft documentation if you need assistance creating a virtual server on the source cluster.

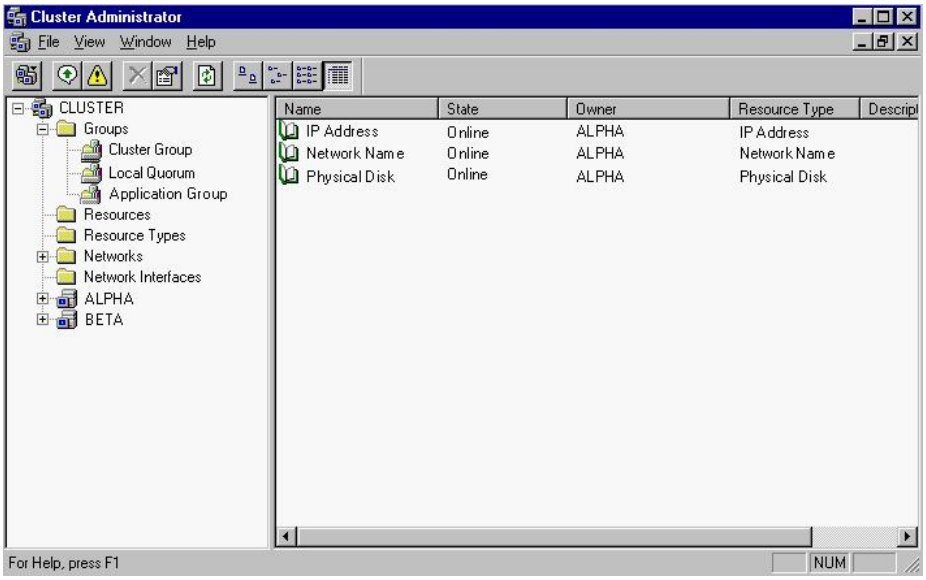

2. If your target is a cluster, create a virtual server (including resources for an IP address, network name, and physical disk) on the target cluster. With this configuration, if there is a source failure, the data will be available for the users from the target cluster, regardless of which node is currently in control. MSCS will handle failover between the nodes of the target cluster. Double-Take Availability will handle failover between the source (cluster or standalone) and the target cluster. See your Microsoft documentation if you need assistance creating a virtual server on the target cluster.
3. On your source, create a replication set from the Replication Console.

Note: If your source is a cluster, you need to create the replication set on the node which currently owns the group with the virtual server you want to protect.

a. Open the Replication Console.
b. Right-click the source in the left pane of the Replication Console and select **New**, **Replication Set**.
c. A replication set icon appears in the left pane under the source. By default, it is named New Replication Set. Rename the newly inserted replication set with a unique name by typing over the default name and pressing **Enter**. This process is similar to naming a new folder in Windows Explorer.
d. Expand the tree under the replication set name to view the volume and directory tree for the source.

Note: The default number of files that are listed in the right pane of the Replication Console is 2500, but this is user configurable. A larger number of file listings allows you to see more files in the Replication Console, but results in a slower display rate. A smaller number of file listings displays faster, but may not show all files contained in the directory. To change the number of files displayed, select **File**, **Options** and adjust the **File Listings** slider bar to the desired number.

To hide offline files, such as those generated by snapshot applications, select **File**, **Options** and disable **Display Offline Files**. Offline files and folders are denoted by the arrow over the lower left corner of the folder or file icon.

e. Identify the data on the source associated with the group that you want to protect by selecting volumes, drives, directories, and/or specific files. Be sure and verify what files can be included in your replication set.

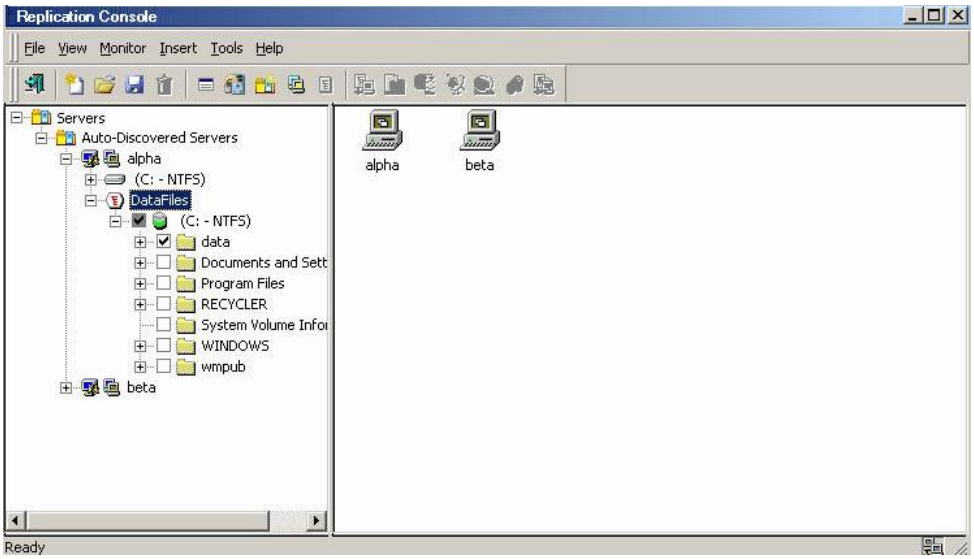

f. After selecting the data for this replication set, right-click the new replication set icon and select **Save**. A saved replication set icon will change from red to black.

4. If your source is a cluster, you need to create a duplicate replication set on each of the other nodes in the cluster. Because the other nodes do not currently own the files, you will not be able to browse to select the data like you did on the first node. Therefore, you will have to manually enter the replication set data.

Note: As an alternative to the following manual steps, you can stop the Double-Take service on the other nodes of the source cluster, copy the file DblTake.db from the first node to the other nodes, and then restart the Double-Take service.

---

a. Right-click the replication set created on the owning node and select **Properties**.

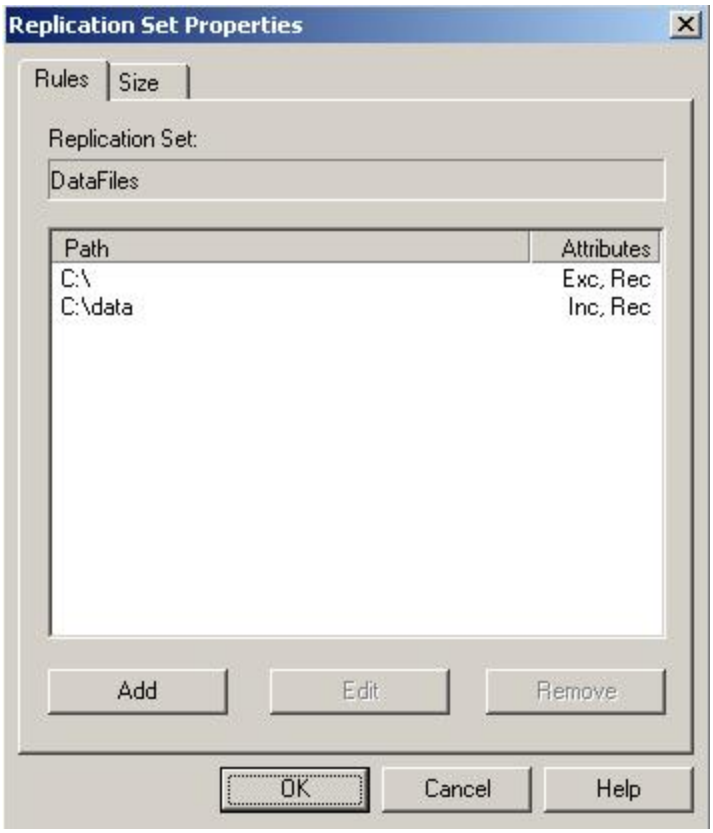

b. Record the exact drive and directories of each path displayed, including where the rule is included or excluded and if recursion is applied.

c. Right-click a non-owning node and select **New**, **Replication Set**.

d. Enter the exact, case-sensitive name for the replication set as specified on the owning node and press **Enter**.

e. Right-click the replication set that you just created and select **Properties**.

f. Click **Add**.

g. Enter the exact same replication set rules you recorded from the owning node. Be sure and mark the correct **Include**, **Exclude** and **Recurse sub-directories** options that need to be applied.

Note: Each replication set rule on the non-owning nodes must be identical to the replication set rule on the owning node.

---

h. After entering all of the replication set rules, save the replication set. The replication set rules will be saved even though the non-owning nodes do not have access to the locations right now. The rules will function properly when the node becomes an owner.

5. On your cluster (source and/or target), you need to disable the standard Double-Take Availability connection controls so that the MSCS resource that you will be configuring later can control the Double-Take Availability connections.
    a. In the Replication Console, right-click a node of the source cluster and select **Properties**.

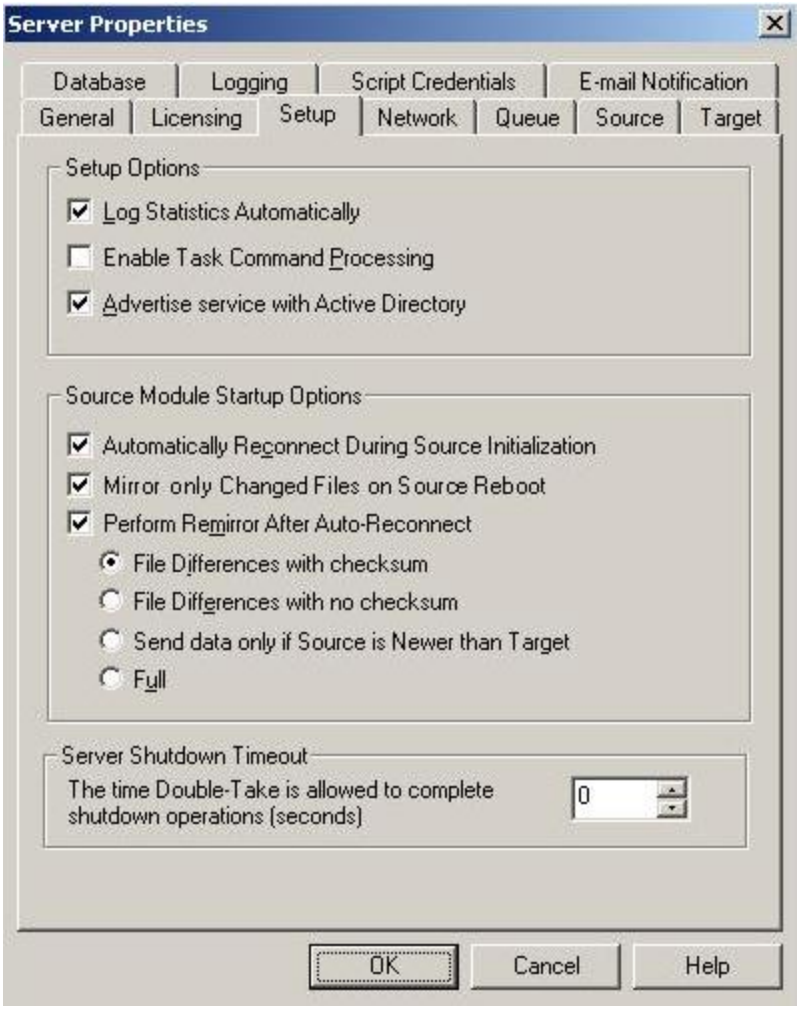

    b. Select the **Setup** tab.
    c. By default, the **Automatically Reconnect During Source Initialization** check box will be selected. Disable this option by clearing the check box.
    d. Click **OK** to save the changes.
    e. Repeat these steps on each node of the cluster(s).
6. If your source is a standalone server, establish your connection through the Connection Manager, selecting the virtual server network name for the **Target Server** and the virtual server IP address for the **Route**.

7. If your source is a cluster, establish your connection by creating and bringing online a Double-Take Source Connection resource. These instructions will vary depending on your operating system.

- Establishing your connection on Windows 2003
- Establishing your connection on Windows 2008

Establishing your connection on Windows 2003

1. From the Cluster Administrator, right-click on the group and virtual server you are protecting and select **New**, **Resource**.
   - **Name**—Specify a name that indicates this is the Double-Take Availability virtual server connection.
   - **Description**—You can optionally add a more detailed description for this resource.
   - **Resource type**—Specify Double-Take Source Connection.
   - **Group**—The resource group name should be selected. If it is not, select the correct group name.
2. Specify the following fields on the New Resource dialog box.
3. Click **Next** to continue.
4. Verify that all of the cluster nodes appear as **Possible Owners** and click **Next** to continue.
5. The resource is dependent on the physical disk and network name resources. Select these two resources so that the Double-Take Source Connection resource is dependent on both of them. Click **Next** to continue.
6. Specify your Double-Take Availability connection parameters.

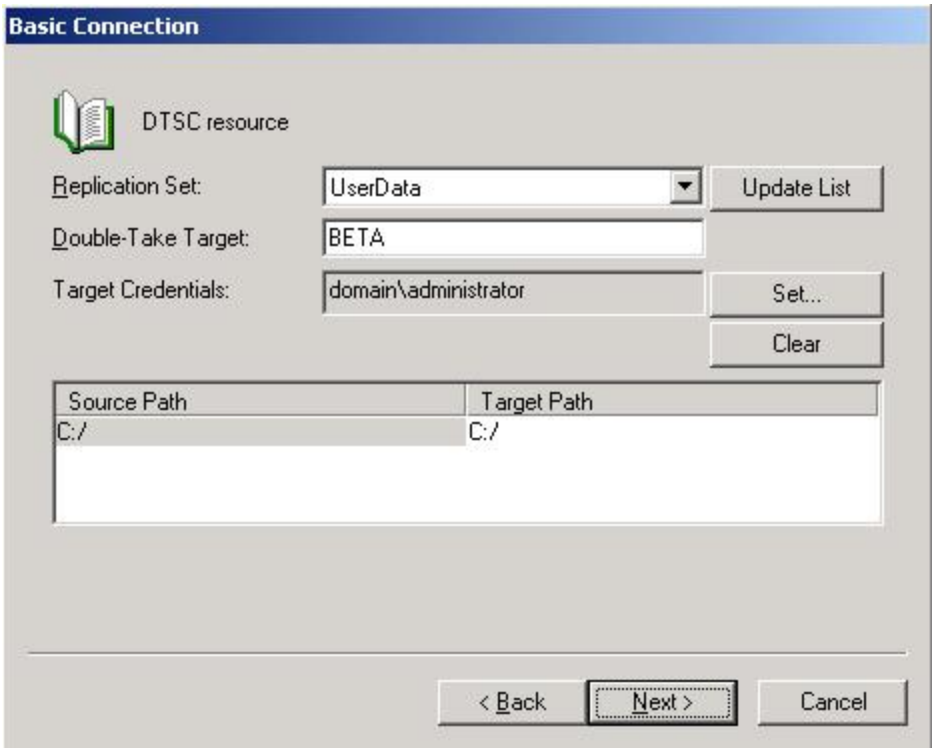

- **Replication Set**—Select the Double-Take Availability replication set that you want to use. If the replication set you want to use is not listed, click **Update List** to refresh the list of replication sets from the source.
- **Double-Take Target**—Specify the name or IP address of the target. If your target is a cluster, this is the virtual name or virtual IP address of the virtual server you created on the target cluster. If your target is a standalone server, this is the name or IP address of the standalone target server.
- **Target Credentials**—To specify the account to use when logging on to the target, click **Set** and enter the user name, password and domain. The user must be a member of the **Double-Take Admin** security group on all nodes of the target cluster or on the standalone target server. Click **OK** to save the settings. If you need to clear the target credentials, click **Clear** then reenter new credentials.
- **Source Path** and **Target Path**—The path(s) for the replication set data from the source is displayed along with the path where the copy of the replication set data will be stored on the target. If you want to change the default target path, click on a path and manually modify the location.

7. Click **Next** to continue.
8. Specify your Double-Take Availability bandwidth limiting parameters.

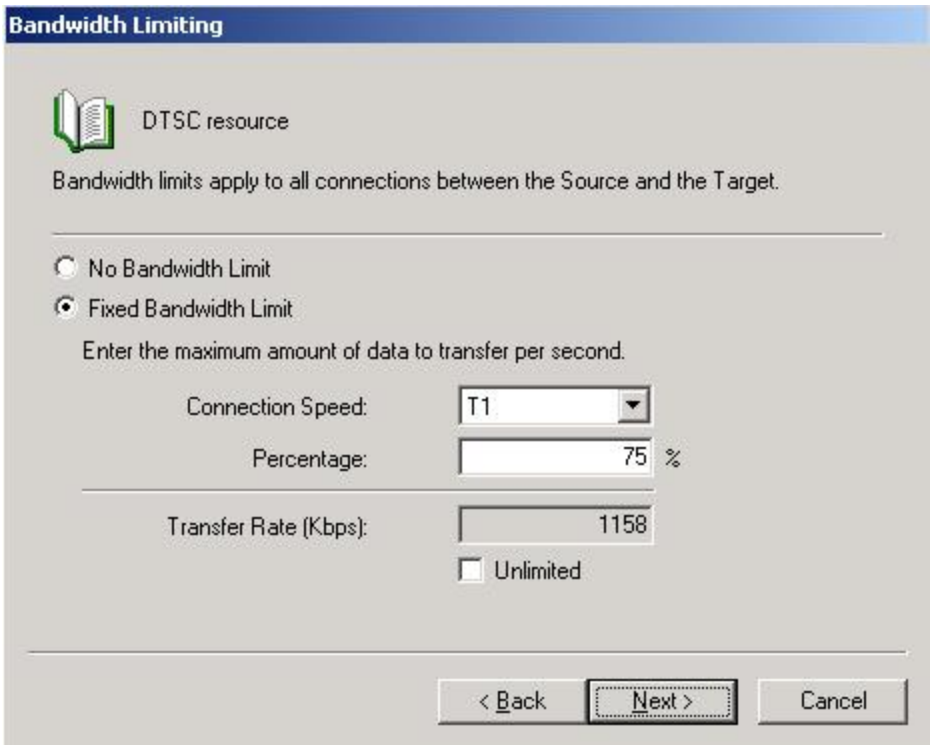

- **No Bandwidth Limit**—Data will be transmitted using all available bandwidth.

- **Fixed Bandwidth Limit**—Data will be transmitted according to the user-specified bandwidth configuration. By default, the **Unlimited** checkbox is enabled. This configuration is identical to selecting **No Bandwidth Limit**. If you want to limit your bandwidth usage, clear this checkbox. To limit the bandwidth usage, enter the maximum amount of data you want to transfer per second. You can indicate it by specifying your **Connection Speed** and the **Percentage** of the bandwidth that you want to use or by entering the **Transfer Rate** value directly.

9. Click **Next** to continue.
10. Specify your Double-Take Availability orphan file parameters. By default, the orphan files feature is disabled. To enable it, select **Move/Delete Orphan Files**. Specify if you want to delete or move the files. If you select the move option, identify the location where these orphan files will be located. Click **Next** to continue.
11. Specify your Double-Take Availability compression parameters. By default, compression is disabled. To enable it, select **Enable Compression**. Depending on the compression algorithms available for your operating system, you may see a slider bar indicating different compression levels. Set the level from minimum to maximum compression to suit your needs. Click **Next** to continue.
12. Specify your Double-Take Availability mirroring settings.

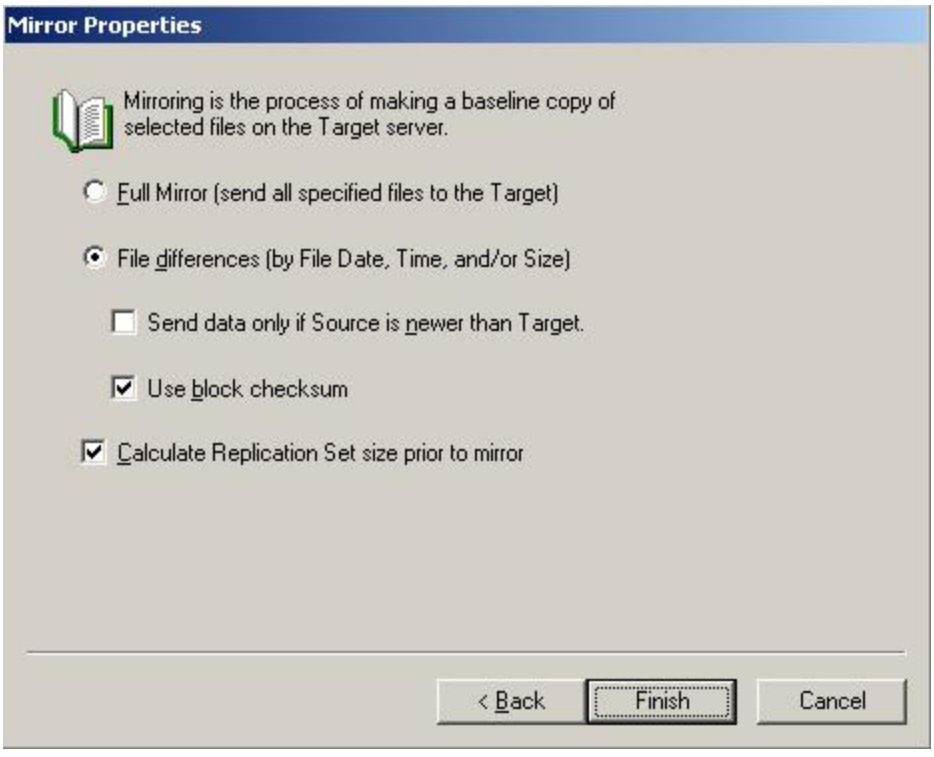

- **Full Mirror**—All files in the replication set will be sent from the source to the target.

- **File differences**—Only those files that are different based size or date and time will be sent from the source to the target.
- **Send data only if Source is newer than Target**—Only those files that are newer on the source are sent to the target.

Note: If you are using a database application, do not use the newer option unless you know for certain you need it. With database applications, it is critical that all files, not just some of them that might be newer, get mirrored.

- **Use block checksum**—For those files flagged as different, the mirror performs a checksum comparison and only sends those blocks that are different.
- **Calculate Replication Set size prior to mirror**—Determines the size of the replication set prior to starting the mirror. The mirroring status will update the percentage complete if the replication set size is calculated.

13. Click **Finish** to complete the creation of the Double-Take Source Connection resource.
14. Bring the Double-Take Source Connection resource and the virtual server resources online.

Establishing your connection on Windows 2008

1. From the Failover Cluster Management applet, right-click on the group and virtual server you are protecting and select **Add a resource**, **More resources**, **Add Double-Take Source Connection**.
2. Right-click on the new resource and select **Properties**.
3. On the **General** tab, specify a **Resource Name** that indicates this is the Double-Take Availability virtual server connection.
4. Select the **Dependencies** tab. This resource is dependent on the physical disk and network name resources. Insert these two resources so that the Double-Take Source Connection resource is dependent on both of them.
5. Select the **Advanced Policies** tab. Verify that all of the cluster nodes appear as **Possible Owners**.
6. Select the **Basic Connection** tab and specify your Double-Take Availability connection parameters.

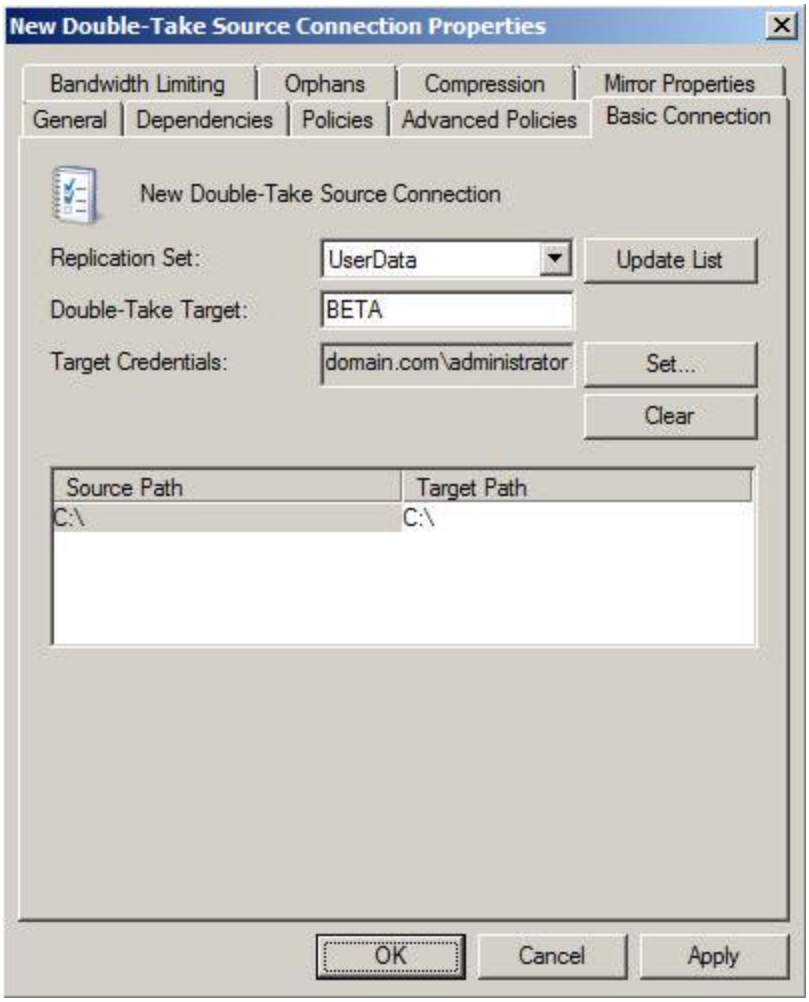

- **Replication Set**—Select the Double-Take Availability replication set that you want to use. If the replication set you want to use is not listed, click **Update List** to refresh the list of replication sets from the source.
- **Double-Take Target**—Specify the name or IP address of the target. If your target is a cluster, this is the virtual name or virtual IP address of the virtual server you created on the target cluster. If your target is a standalone server, this is the name or IP address of the standalone target server.
- **Target Credentials**—To specify the account to use when logging on to the target, click **Set** and enter the user name, password and domain. The user must be a member of the **Double-Take Admin** security group on all nodes of the target cluster or on the standalone target server. Click **OK** to save the settings. If you need to clear the target credentials, click **Clear** then reenter new credentials.
- **Source Path** and **Target Path**—The path(s) for the replication set data from the source is displayed along with the path where the copy of the replication set data will be stored on the target. If you want to change the default target path, click on a path and manually modify the location.

7. Select the **Bandwidth Limiting** tab and specify your Double-Take Availability bandwidth limiting parameters.

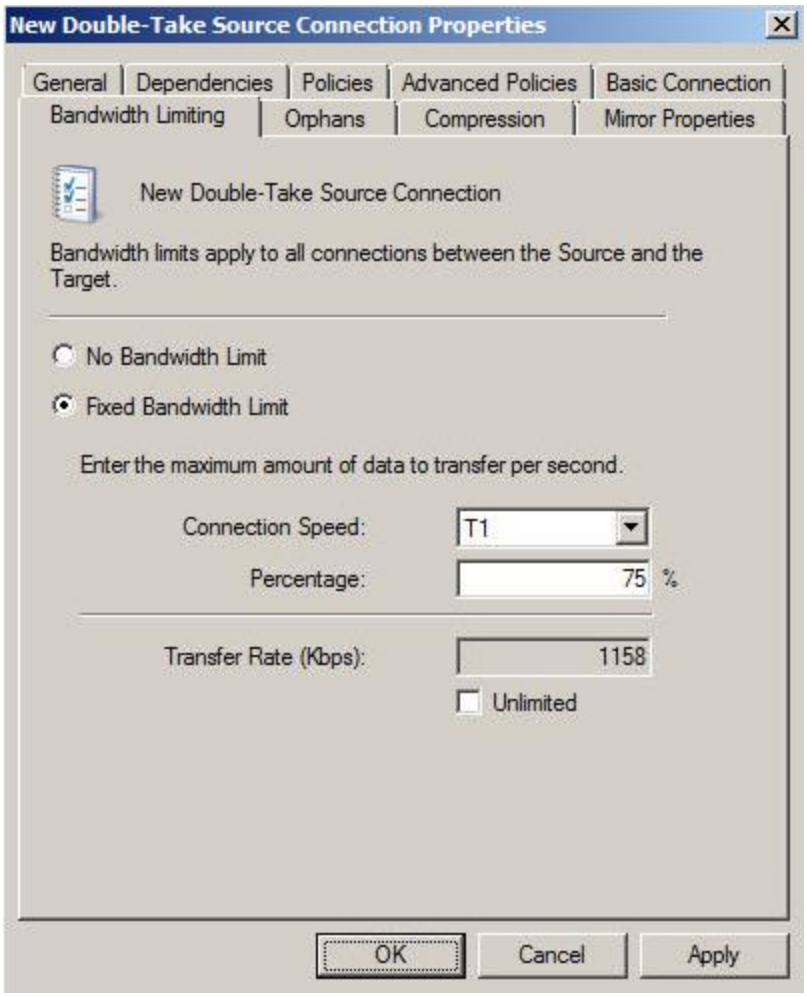

- **No Bandwidth Limit**—Data will be transmitted using all available bandwidth.
- **Fixed Bandwidth Limit**—Data will be transmitted according to the user-specified bandwidth configuration. By default, the **Unlimited** checkbox is enabled. This configuration is identical to selecting **No Bandwidth Limit**. If you want to limit your bandwidth usage, clear this checkbox. To limit the bandwidth usage, enter the maximum amount of data you want to transfer per second. You can indicate it by specifying your **Connection Speed** and the **Percentage** of the bandwidth that you want to use or by entering the **Transfer Rate** value directly.

8. Select the **Orphans** tab and specify your Double-Take Availability orphan file parameters. By default, the orphan files feature is disabled. To enable it, select **Move/Delete Orphan Files**. Specify if you want to delete or move the files. If you select the move option, identify the location where these orphan files will be located.

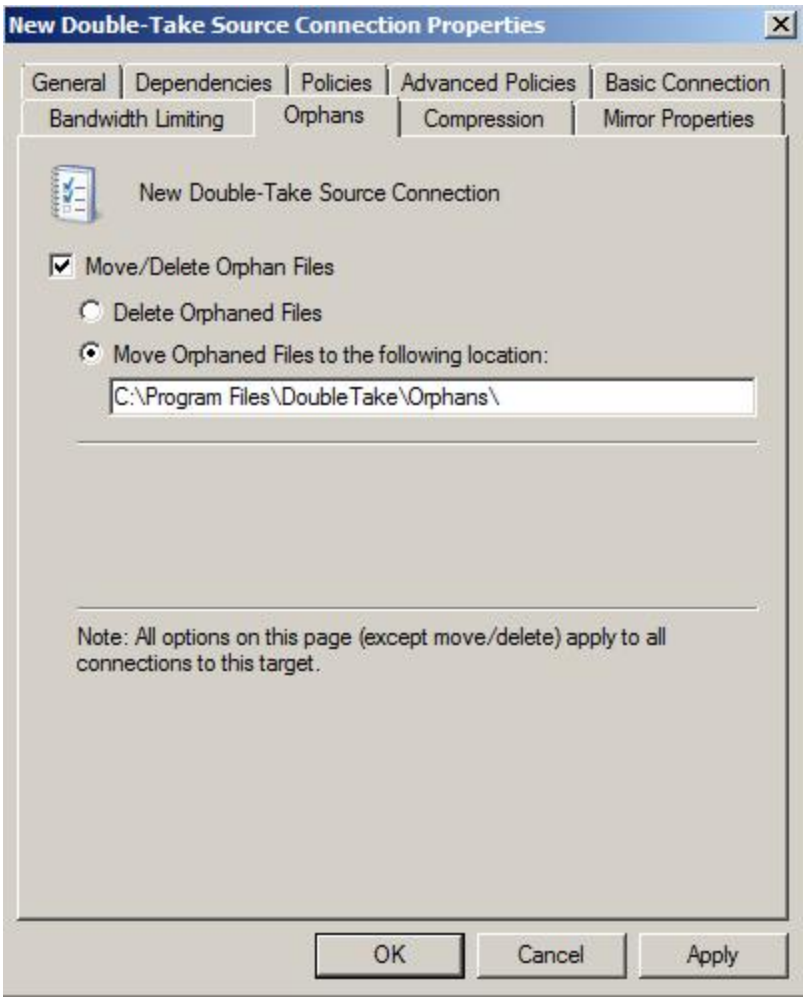

9. Select the **Compression** tab and specify your Double-Take Availability compression parameters. By default, compression is disabled. To enable it, select **Enable Compression**. Depending on the compression algorithms available for your operating system, you may see a slider bar indicating different compression levels. Set the level from minimum to maximum compression to suit your needs.

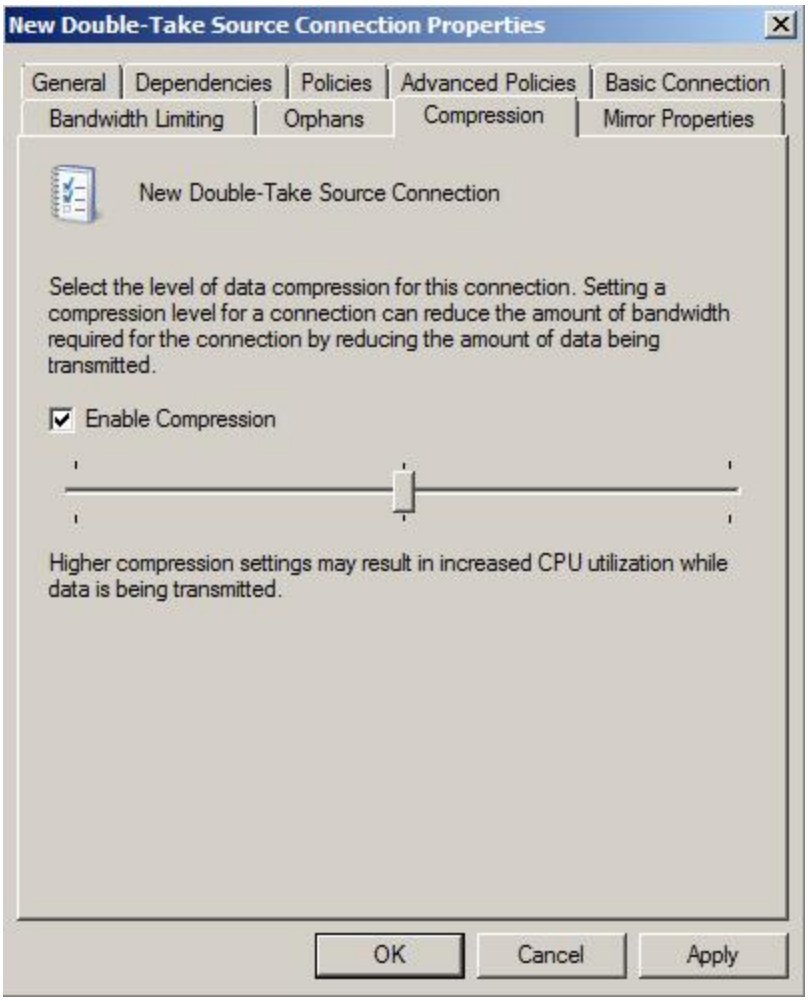

10. Select the **Mirror Properties** tab and specify how you want to mirror the data.

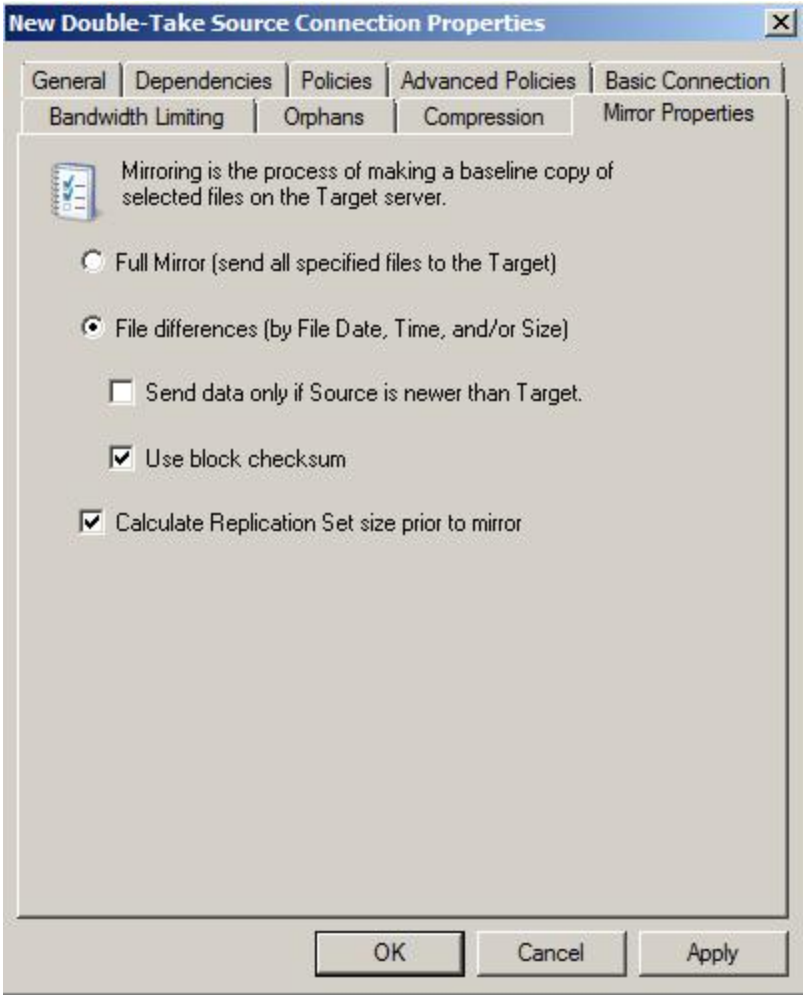

- **Full Mirror**—All files in the replication set will be sent from the source to the target.
- **File differences**—Only those files that are different based size or date and time will be sent from the source to the target.
- **Send data only if Source is newer than Target**—Only those files that are newer on the source are sent to the target.

Note: If you are using a database application, do not use the newer option unless you know for certain you need it. With database applications, it is critical that all files, not just some of them that might be newer, get mirrored.

- **Use block checksum**—For those files flagged as different, the mirror performs a checksum comparison and only sends those blocks that are different.
- **Calculate Replication Set size prior to mirror**—Determines the size of the replication set prior to starting the mirror. The mirroring status will update the percentage complete if the replication set size is calculated.

11. Click **OK** to complete the creation of the Double-Take Source Connection resource.
12. Bring the Double-Take Source Connection resource and the virtual server resources online.

Protecting a GeoCluster

The GeoCluster Replicated Disk resource allows for the real-time copy of data to be available on other nodes in the cluster. In the event of a failure and another node takes ownership, the GeoCluster Replicated Disk resource is also moved to the other node and it continues to replicate data, in real-time, to the remaining nodes in the cluster.

The instructions for creating this resource are different depending on your operating system.

- Creating the GeoCluster Replicated Disk Resource on Windows 2003
- Creating the GeoCluster Replicated Disk Resource on Windows 2008
- Bringing the resource online
- Taking the resource offline

Creating the GeoCluster Replicated Disk Resource on Windows 2003

1. Select **Start**, **Programs**, **Administrative Tools**, **Cluster Administrator**.
2. Right-click the group that you want to add a replicated disk to and select **New**, **Resource**.
3. Specify the following fields on the New Resource dialog box.
    - **Name**—Specify a name that identifies which application, file set, disk, and so on that you are protecting. This name must be unique within the cluster.
    - **Description**—You can optionally add a more detailed description for this resource.
    - **Resource type**—Specify GeoCluster Replicated Disk.
    - **Group**—The group that you originated the new resource from will be selected. Verify that this is the correct group. If it is not, select the correct group name.
4. Click **Next** to continue.
5. The GeoCluster Replicated Disk resource ensures that an up-to-the-minute copy of the data resides on all nodes identified in the **Possible owners** list. All nodes are included in the default, which should not be changed. Click **Next** to continue.
6. The GeoCluster Replicated Disk resource is not dependent on any other resources. Click **Next** to continue.
7. Specify the GeoCluster Replicated Disk parameters using the settings below.
    - **Disk to replicate**—Select a disk to replicate from the available volumes. The only volumes that will be displayed are those that meet the following criteria.
        - NTFS volumes
        - Volumes which are not already being replicated by another GeoCluster Replicated Disk resource
        - Volumes that are not physical disk resources
        - Volumes that do not contain system files (The volume that you booted Windows from will not be displayed.)
        - Volumes that exist on all nodes of the cluster
    - **Network**—Select the network that you want to use for Double-Take Availability mirroring and replication traffic. If you do not have multiple networks established, you will only be able to select the one network that does exist. If you do not select a network, Double-Take Availability will use DNS to determine a network route to use. Ideally, the networks used for various traffic should be separated. This is dependent on the number of networks that you established when you created the cluster and the priority assigned to each network. For example, if you have two network routes,

separate Double-Take Availability and your public traffic. If you have three routes, separate the public traffic and then separate Double-Take Availability from the cluster heartbeat.

- **Orphan files**—An orphan is a file that exists in the target location but is not in the source location. You can enable the resource to remove or delete orphans during a mirror. If you choose to move orphan files, specify the location where you want to move them. You should not select the system volume for the orphan files because it could impact the stability of the cluster service.

8. Click **Next** to continue.
9. If you want to configure compression, verify that **Enable Compression** is selected. Depending on the compression algorithms available for your operating system, you may see a slider bar indicating different compression levels. Set the level from minimum to maximum compression to suit your needs.
10. Click **Next** to continue.
11. Specify your Double-Take Availability mirroring settings.
    - **Full Mirror**—All files in the replication set will be sent from the source to the target.
    - **File differences**—Only those files that are different based size or date and time will be sent from the source to the target.
      - **Send data only if Source is newer than Target**—Only those files that are newer on the source are sent to the target.

        Note: If you are using a database application, do not use the newer option unless you know for certain you need it. With database applications, it is critical that all files, not just some of them that might be newer, get mirrored.

      - **Use block checksum**—For those files flagged as different, the mirror performs a checksum comparison and only sends those blocks that are different.
    - **Calculate Replication Set size prior to mirror**—Determines the size of the replication set prior to starting the mirror. The mirroring status will update the percentage complete if the replication set size is calculated.
12. Click **Finish** to complete the creation of the GeoCluster Replicated Disk resource.

Creating the GeoCluster Replicated Disk Resource on Windows 2008

1. Select **Start**, **Programs**, **Administrative Tools**, **Failover Cluster Management**.
2. Right-click the application group where you want to add a replicated disk to and select **Add a resource**, **More resources**, **Add GeoCluster Replicated Disk**.
3. Right-click on the resource and select **Properties**.
4. On the **Connection Parameters** tab, specify the GeoCluster Replicated Disk connection parameters using the settings below.
    - **Disk to replicate**—Select a disk to replicate from the available volumes. The only volumes that will be displayed are those that meet the following criteria.
        - NTFS volumes
        - Volumes which are not already being replicated by another GeoCluster Replicated Disk resource
        - Volumes that are not physical disk resources
        - Volumes that do not contain system files (The volume that you booted Windows from will not be displayed.)
        - Volumes that exist on all nodes of the cluster

        Note: If you are using Hyper-V, select the drive where the virtual machine .vhd file is stored.

    - **Network to route Double-Take mirroring and replication traffic over**—Select the network that you want to use for Double-Take Availability mirroring and replication traffic. If you do not have multiple networks established, you will only be able to select the one network that does exist. If you do not select a network, Double-Take Availability will use DNS to determine a network route to use. Ideally, the networks used for various traffic should be separated. This is dependent on the number of networks that you established when you created the cluster and the priority assigned to each network. For example, if you have two network routes, separate Double-Take Availability and your public traffic. If you have three routes, separate the public traffic and then separate Double-Take Availability from the cluster heartbeat.
    - **Interval to check unresponsive nodes**—Specify how much time, in seconds, between checks of nodes to see if a Double-Take Availability connection can be made.
    - **Delay connection until resources dependent on this one are online**—This option allows you to delay a Double-Take Availability connection until

any resources that have the GeoCluster Replicated Disk resource as a dependency are online. By ensuring that all resources that are dependent on the GeoCluster Replicated disk resource are online before starting the connection, the chance of a conflict occurring because application resources are attempting to open files exclusively while Double-Take Availability is mirroring those files is removed.

- **Enable orphans**—An orphan is a file that exists in the target location but is not in the source location. You can enable the resource to Move orphans or Delete orphans during a mirror. If you choose to move orphan files, specify the location where you want to move them. You should not select the system volume for the orphan files because it could impact the stability of the cluster service.

5. No other settings are required for the GeoCluster Replicated Disk resource, although there are optional settings available. Click **OK** to save the GeoCluster Replicated Disk configuration changes that you made.
6. If you are using Hyper-V, you have several additional configuration steps.
   a. Take the virtual machine's configuration resource offline by right-clicking on it and selecting **Take this resource offline**.
   b. Modify the virtual machine's configuration resource and specify the GeoCluster Replicated Disk resource (not the GRD Status resource) as a dependency.
   c. Modify the virtual machine resource (the parent of the virtual machine configuration resource) and specify the GeoCluster Replicated Disk resource (not the GRD Status resource) as a dependency.
   d. From a command prompt, execute the following command, substituting the name of your GeoCluster Replicated Disk resource.

      cluster res "New Geocluster Replicated Disk" /priv sourcefailoverdelay=15

      The GeoCluster Replicated Disk will mirror and replicate the virtual machine data to the other cluster nodes.
7. To control the resource, you can bring it online and take it offline. Neither of these actions trigger failover. They just control the activity of the resource.

Bringing the resource online

The GeoCluster Replicated Disk resource will appear offline after it is created. When you bring it online, the following actions occur.

1. A Double-Take Availability replication set is created with the same name that was assigned to the resource.
2. The replication set is connected to all of the possible owners specified in the resource.
3. A mirror is initiated to create the baseline copy of data from the active node to all of the possible owners.
4. The drive where the mirrored data is located on each of the possible owners is made read-only to all other applications except Double-Take Availability.
5. Real-time replication from the active node to all of the possible owners begins.

If you are using Windows 2003, right-click the resource and select **Bring online**.

If you are using Windows 2008, right-click the resource and select **Bring this resource online**.

Taking the resource offline

When you take the GeoCluster Replicated Disk resource offline, the following actions occur.

1. Real-time replication from the active node to the possible owners stops.
2. The read-only limitation is removed from the corresponding drive letters on the possible owners.
3. The replication set is disconnected from all of the possible owners.
4. The replication set is deleted.

If you are using Windows 2003, right-click the resource and select **Take offline**.

If you are using Windows 2008, right-click the resource and select **Take this resource offline**.

Note: If the GeoCluster Replicated Disk Resource is offline, data integrity cannot be guaranteed on the other nodes in the cluster.

GeoCluster resource properties

Resource properties are displayed differently in Windows 2003 and Windows 2008. For example, the possible owners of a resource is listed on the **General** tab of the resource properties in Windows 2003, while in Windows 2008 they are listed on the **Advanced Policies** tab.

For both operating systems, right-click the resource and select **Properties**, when you want to view or modify the resource properties.

- GeoCluster Replicate Disk properties on Windows 2003
- GeoCluster Replicate Disk properties on Windows 2008

GeoCluster Replicated Disk properties on Windows 2003

There are six properties tabs for the GeoCluster Replicated Disk resource on Windows 2003.

1. **General**—This tab identifies the **Name** and **Description** of the resource and the **Possible Owners**. If you change the name of the resource, the replication set name will not change until the next time the resource is brought online. The GeoCluster Replicated Disk resource must have at least two possible owners to function properly. Modifying the **Possible Owners** list will cause one of the following actions to occur.
    - If you add additional **Possible Owners**, the GeoCluster Replicated Disk resource will connect the resource's replication set to the new owners and begin a mirror to each.
    - If you remove **Possible Owners**, the GeoCluster Replicated Disk resource will disconnect the resource's replication set from each owner removed.
2. **Dependencies**—By default, the GeoCluster Replicated Disk resource is not dependent on any other resources.
3. **Advanced settings**—This tab controls how and when MSCS handles a failure of the resource.
    - **Do not restart**—Select this option if you do not want cluster service to restart the resource if it fails.
    - **Restart**—Select this option if you want cluster service to restart the resource if it fails.
        - Enable **Affect the group** if you want a failure of this resource to move the group to another node. If you disable this option, cluster service still attempts to restart the resource using the **Threshold** and **Period** values, but the failure of the resource will not cause the group to move to another node.
        - The **Threshold** and **Period** values determine the number of times cluster service will attempt to restart the failed resource within a specified period of time before moving the group to another node.
    - **"Looks Alive" poll interval**—This setting specifies how often the resource is polled to determine whether it is still running on the active node. You can choose a value from the resource type, or you can specify your own value.
    - **"Is Alive" poll interval**—This setting designates how often the possible owners are polled to determine whether the specified disk on each node can be written to and read from. You can choose a value from the resource type, or you can specify your own value.
    - **Pending timeout**—This value determines how long the resource is allowed

to remain in a pending state before it fails. If the resource takes longer than the time specified to be brought online or taken offline, the resource fails.

For more information on the **Advanced Setting**s options, see your Windows documentation.

4. **Connection parameters**—This tab controls disk replication, network routing, and orphan files for GeoCluster.
   - **Disk to replicate**—The volume to replicate
   - **Network to route Double-Take mirroring and replication traffic over**—The network to use for Double-Take Availability mirroring and replication traffic. If you do not have multiple networks established, you will only be able to select the one network that does exist. If you do not select a network, GeoCluster will use DNS to determine a network route to use.

     Ideally, the networks used for various traffic should be separated. This is dependent on the number of networks that you established when you created the cluster and the priority assigned to each network. For example, if you have two network routes, separate GeoCluster and your public traffic. If you have three routes, separate the public traffic and then separate GeoCluster from the cluster heartbeat.

     Modifications to either of the first two settings will not take effect until the next time the resource is brought online.
   - **Interval to check unresponsive nodes**—The frequency to determine how often an unresponsive node is checked to see if a Double-Take Availability connection can be made
   - **Delay connection until resources dependent on this one are online**—This option allows you to delay a Double-Take Availability connection until any resources that have the GeoCluster Replicated Disk resource as a dependency are online. By ensuring that all resources that are dependent on the GeoCluster Replicated disk resource are online before starting the connection, the chance of a conflict occurring because application resources are attempting to open files exclusively while GeoCluster is mirroring those files is removed.
   - **Orphan files**—An orphan is a file that exists in the target location but is not in the source location. You can enable the resource to remove or delete orphans during a mirror. If you choose to move orphan files, specify the location where you want to move them. You should not select the system volume for the orphan files because it could impact the stability of the cluster service.
5. **Compression**—If you want to configure Double-Take Availability compression, verify that **Enable Compression** is selected. Depending on the compression

algorithms available for your operating system, you may see a slider bar indicating different compression levels. Set the level from minimum to maximum compression to suit your needs.

6. **Mirror Properties**—This tab controls the Double-Take Availability mirroring process.
    - **Full Mirror**—All files in the replication set will be sent from the source to the target.
    - **File differences**—Only those files that are different based size or date and time will be sent from the source to the target.
        - **Send data only if Source is newer than Target**—Only those files that are newer on the source are sent to the target.

          Note: If you are using a database application, do not use the newer option unless you know for certain you need it. With database applications, it is critical that all files, not just some of them that might be newer, get mirrored.

        - **Use block checksum**—For those files flagged as different, the mirror performs a checksum comparison and only sends those blocks that are different.
    - **Calculate Replication Set size prior to mirror**—Determines the size of the replication set prior to starting the mirror. The mirroring status will update the percentage complete if the replication set size is calculated.

GeoCluster Replicated Disk properties on Windows 2008

There are eight properties tabs for the GeoCluster Replicated Disk resource on Windows 2008.

1. **General**—This tab identifies the **Name** and **Resource type** of the resource. It also displays the current state of the resource and an additional detailed status message.
2. **Dependencies**—By default, the GeoCluster Replicated Disk resource is not dependent on any other resources.
3. **Policies**—This tab controls how and when MSCS handles a failure of the resource. For more information on **Policies** options, see your Windows documentation.
    - **If resource fails, do no restart**—Select this option if you do not want cluster service to restart the resource if it fails.
    - **If resource fails, attempt restart on current node**—Select this option if you want cluster service to restart the resource if it fails. Specify the length of time to attempt restarts and the number of restarts to attempt during that period of time.
    - **If restart is unsuccessful, fail over all resources in this service or application**—If this option is enabled, the failure of the group will cause the resource to move to another node. If this option is disabled, the failure of the resource will not cause the resource to move to another node.
    - **If all the restart attempts fail, begin restarting again after the specified period**—If this option is enabled, the cluster will delay the length of time specified before trying to restart the resource again.
    - **Pending timeout**—This value determines how long the resource is allowed to remain in a pending state before it fails. If the resource takes longer than the time specified to be brought online or taken offline, the resource fails.
4. **Advanced Policies**—This tab controls resource specific settings. For more information on **Advanced Policies** options, see your Windows documentation.
    - **Possible owners**—All nodes of the cluster are listed. Select or deselect the nodes that you want to be possible owners.
        - If you add additional owners, the GeoCluster Replicated Disk resource will connect the resource's replication set to the new owners and begin a mirror to each.
        - If you remove owners, the GeoCluster Replicated Disk resource will disconnect the resource's replication set from each owner removed.

The GeoCluster Replicated Disk resource must have at least two possible owners to function properly.

- **Basic resource health check interval**—This setting is formerly known as the Looks Alive poll interval. It specifies how often the resource is polled to determine whether it is still running on the active node. You can choose the standard time period of 5 seconds, or you can specify your own value.
- **Thorough resource health check interval**—This setting is formerly known as the Is Alive poll interval. It designates how often the possible owners are polled to determine whether the specified disk on each node can be written to and read from. You can choose the standard time period of 1 minute, or you can specify your own value.
- **Run this resource in a separate Resource Monitor**—You should enable this option so that each GeoCluster Replicated Disk resource runs in its own monitor.

5. **Connection parameters**—This tab controls disk replication, network routing, and orphan files for Double-Take Availability.
   - **Disk to replicate**—The volume to replicate
   - **Network to route Double-Take mirroring and replication traffic over**—The network to use for Double-Take Availability mirroring and replication traffic. If you do not have multiple networks established, you will only be able to select the one network that does exist. If you do not select a network, GeoCluster will use DNS to determine a network route to use.

     Ideally, the networks used for various traffic should be separated. This is dependent on the number of networks that you established when you created the cluster and the priority assigned to each network. For example, if you have two network routes, separate GeoCluster and your public traffic. If you have three routes, separate the public traffic and then separate GeoCluster from the cluster heartbeat.

     Modifications to either of the first two settings will not take effect until the next time the resource is brought online.
   - **Interval to check unresponsive nodes**—The frequency to determine how often an unresponsive node is checked to see if a Double-Take Availability connection can be made
   - **Delay connection until resources dependent on this one are online**—This option allows you to delay a Double-Take Availability connection until any resources that have the GeoCluster Replicated Disk resource as a dependency are online. By ensuring that all resources that are dependent on the GeoCluster Replicated disk resource are online before starting the connection, the chance of a conflict occurring because application resources

are attempting to open files exclusively while GeoCluster is mirroring those files is removed.

- **Enable orphans**—An orphan is a file that exists in the target location but is not in the source location. You can enable the resource to remove or delete orphans during a mirror. If you choose to move orphan files, specify the location where you want to move them. You should not select the system volume for the orphan files because it could impact the stability of the cluster service.

6. **Compression**—If you want to configure Double-Take Availability compression, verify that **Enable Compression** is selected. Depending on the compression algorithms available for your operating system, you may see a slider bar indicating different compression levels. Set the level from minimum to maximum compression to suit your needs.
7. **Online Pending**—Because context-sensitive, right-click menus are not available in the Windows 2008 Failover Cluster Administrator, GeoCluster processing controls have been added to a properties tab. For details on this tab, see Monitoring a cluster workload.
8. **Mirror Properties**—This tab controls the Double-Take Availability mirroring process.
    - **Full Mirror**—All files in the replication set will be sent from the source to the target.
    - **File differences**—Only those files that are different based size or date and time will be sent from the source to the target.
        - **Send data only if Source is newer than Target**—Only those files that are newer on the source are sent to the target.

          Note: If you are using a database application, do not use the newer option unless you know for certain you need it. With database applications, it is critical that all files, not just some of them that might be newer, get mirrored.

        - **Use block checksum**—For those files flagged as different, the mirror performs a checksum comparison and only sends those blocks that are different.
    - **Calculate Replication Set size prior to mirror**—Determines the size of the replication set prior to starting the mirror. The mirroring status will update the percentage complete if the replication set size is calculated.

Configuring failover monitoring

If you are configuring failover monitoring for a cluster workload, use the same process for a data workload, except note the following changes.

- If your source is a cluster, you must use the **Custom** option when identifying the source machine and specify the virtual network name and virtual IP address on the source cluster.
- You must use a pre-failback and post-failover script. The pre-failback script must contain the following case-sensitive command, substituting the name of the IP address resource assigned to the target virtual server.

  Cluster resource “IP_Address_Resource_Name” /OFFLINE
- The post-failback script must contain the following case-sensitive command, substituting the name of the group on the target virtual server.

  Cluster group “Group_Name” /ONLINE
- If your target is a cluster, you need to repeat the steps to establish failover monitoring on all nodes of the cluster, so that any node that could be the owning node of the target cluster will be monitoring the source.

Special configurations

Double-Take Availability can be implemented with very little configuration necessary in small or simple networks, but additional configuration may be required in large or complex environments. Because an infinite number of network configurations and environments exist, it is difficult to identify all of the possible configurations. Select a link to review configuration information for particular network environments.

- Domain controllers
- NetBIOS
- WINS
- DNS
- Non-Microsoft DNS
- Macintosh shares
- NFS Shares

Domain controllers

Failover of domain controllers is dependent on the Double-Take Availability functionality you are using.

- **Domain controller role**—If you want to failover a domain controller, including the roles of the domain controller, you should create full-server or virtual workload protection.
- **Non-domain controller role**—If you are only protecting data, you can failover a domain controller, but the roles of the domain controller are not included. The server will be a member server after failover. In this case, you need to keep in mind that the unavailability of some of the FSMO (Flexible Single Master Operation) roles can cause immediate issues such as the inability to extend the Active Directory schema or to add a domain to a forest.
- **Global catalog server**—If your source is a global catalog server, you should have other global catalog servers throughout the network to ensure that the failure of the source will not impact users.

NetBIOS

Because NetBIOS is not available on Windows 2008, if you are using Windows 2008 in a workgroup environment and do not have DNS host records, there may be a delay of up to five minutes before the failed over source server name is available for client access. See *About the NetBIOS Interface* in the MSDN library.

WINS

When Double-Take Availability failover occurs, Windows initiates WINS registration with the target's primary WINS server to associate the source server's name with the target's primary IP address. In an environment with just one WINS server, no additional processing is required. In an environment with more than one WINS server, WINS replication will distribute the updated WINS registration to other WINS servers on the network. The length of time required for all WINS servers to obtain the new registration depends on the number of WINS servers, the WINS replication architecture, and the WINS replication interval. Clients will be unable to access the target until their WINS server has received the updated WINS information. You can reduce the time required for the WINS updates, thereby decreasing the wait time for the end users, by scripting the WINS updates in the Double-Take Availabilityfailover scripts. You have two options for scripting the WINS updates.

- WINS registration—This option registers a user-specified server with WINS. It requires less network overhead but administrator group membership on all WINS servers.
- WINS replication—This option forces WINS replication. It does not require any special privileges, but requires system and network resources to complete WINS replication. The impact on the network will depend on the size and complexity of the WINS architecture.

WINS registration

WINS registration can be added to your failover and failback scripts by using the Windows NETSH command with the WINS add name context. Add the following command to your failover and failback scripts for each additional WINS server in your environment (excluding the target's primary WINS server).

netsh wins server ***wins_server_IP_address*** add name Name=***source_server_name*** RecType=1 IP={***IP_address***}

Use the following variable substitutions.

- ***wins_server_IP_address***—The IP address of the WINS server
- ***source_server_name***—The name of the source server
- ***IP_address***—The IP address of the target that has taken over for the failed source (for the failover script) or the IP address of the source that is reassuming its original identity (for the failback script)

For example, suppose you had the following environment.

- **Source name and IP address**—Alpha 192.168.1.108
- **Target name and IP address**—Beta 116.123.2.47
- **Target's Primary WINS server**—116.123.2.50
- **First secondary WINS server on the network**—192.168.1.110
- **Second secondary WINS server on the network**—150.172.114.74

You would add the following to your failover script to register the source's name with the target's IP address on the two secondary WINS servers.

```
netsh wins server 192.168.1.110 add name Name=Alpha RecType=1
IP={116.123.2.47}
netsh wins server 150.172.114.74 add name Name=Alpha RecType=1
IP={116.123.2.47}
```

You would add the following to your failback script to register the source's name back with the source's original IP address on the two secondary WINS servers.

```
netsh wins server 192.168.1.110 add name Name=Alpha RecType=1
IP={192.168.1.108}
netsh wins server 150.172.114.74 add name Name=Alpha RecType=1
IP={192.168.1.108}
```

See your Windows documentation or the Microsoft web site for more details on the NETSH command.

WINS replication

WINS replication can be added to your failover and failback scripts by using the Windows NETSH command with the WINS set replicate context. Add the following command to your failover and failback scripts.

netsh wins server ***target's_primary_wins_server_IP_address*** set replicateflag 1

Use the following variable substitution.

- ***target's_primary_wins_server_IP_address***—The IP address of the target's primary WINS server

For example, suppose you had the following environment.

- **Source name and IP address**—Alpha 192.168.1.108
- **Target name and IP address**—Beta 116.123.2.47
- **Target's Primary WINS server**—116.123.2.50
- **First secondary WINS server on the network**—192.168.1.110
- **Second secondary WINS server on the network**—150.172.114.74

You would add the following to your failover script to force the target's primary WINS server to replicate its updated information to the other secondary WINS servers on the network.

```
netsh wins server 116.123.2.50 set replicateflag 1
```

You would add the same line to your failback script to force the target's primary WINS server to replicate its updated information again. This would replicate information for the source's name and the source's original IP address to the other secondary WINS servers on the network.

```
netsh wins server 116.123.2.50 set replicateflag 1
```

See your Windows documentation or the Microsoft web site for more details on the NETSH command.

DNS

If you are using a Microsoft DNS server, when Double-Take Availability failover occurs, DNS is not automatically updated. If the end-users use DNS to resolve server names and the source IP address was not failed over to the target, additional DNS updates will be required because the host records for the source will remain intact after failover. You can automate this process by scripting the DNS updates in the failover and failback scripts. You have two options for scripting the DNS updates.

- Windows DNSCMD command—The Windows Support Tools contain a DNS Server Troubleshooting Tool utility. This utility includes the DNSCMD command which can be scripted to delete and add host and reverse lookup entries in DNS.
- Double-Take Availability DFO utility—Double-Take Availability also has a utility, called DFO (DNS Failover). The DFO utility can be used to script the deletion and addition of the host and reverse lookup entries in DNS. This utility can be found on the product CD or from the Double-Take Software support web site.

Windows DNSCMD command

DNS updates can be added to your failover and failback scripts by using the Windows DNSCMD command as long as dynamic updates are enabled on the DNS zone and the account running the Double-Take service is a member of the DNSAdmins security group. (See your Microsoft documentation to verify if dynamic updates are enabled.) Add the following commands to your failover and failback scripts to delete the host and reverse lookup entries and add new entries associating the source to the target.

- dnscmd ***DNS_server's_FQDN*** /RecordDelete ***DNS_zone source_server_name*** A ***source_server_IP_address*** /f
- dnscmd ***DNS_server's_FQDN*** /RecordDelete ***www.xxx***.in-addr.arpa ***zzz.yyy*** PTR ***source_server's_FQDN*** /f
- dnscmd ***DNS_server's_FQDN*** /RecordAdd ***DNS_zone source_server_name*** A ***target_server_IP_address***
- dnscmd ***DNS_server's_FQDN*** /RecordAdd ***aaa.bbb***.in-addr.arpa ***ddd.ccc*** PTR ***source_server's_FQDN***

Use the following variable substitutions.

- ***DNS_server's_FQDN***— The fully qualified domain name of the DNS server
- ***DNS_zone*** —The name of the DNS zone
- ***source_server_name*** —The name of the source server
- ***source_server_IP_address***—The IP address on the source
- ***www.xxx***—The first two octets of the source's IP address. For example, if the source's IP address is 192.168.1.108, this variable would be 192.168.
- ***zzz.yyy***—The last two octets, in reverse order, of the source's IP address. For example, if the source's IP address is 192.168.1.108, this variable would be 108.1.
- ***source_server's_FQDN***—The fully qualified domain name of the source server
- ***target_server_IP_address***—The IP address on the target
- ***aaa.bbb***—The first two octets of the target's IP address. For example, if the target's IP address is 116.123.2.47, this variable would be 116.123.
- ***ddd.ccc***— The last two octets, in reverse order, of the target's IP address. For example, if the target's IP address is 116.123.2.47, this variable would be 47.2.

For example, suppose you had the following environment.

- **Full qualified domain name of the source**—Alpha.domain.com
- **Source IP address**—192.168.1.108
- **Fully qualified domain name of the target**—Beta.domain.com
- **Target IP address**—116.123.2.47

- **Fully qualified domain name of the DNS server**—DNSServer.domain.com
- **DNS zone**—domain.com

You would add the following to your failover script to delete the host and reverse lookup entries and add new entries associating the source to the target.

```
dnscmd DNSServer.domain.com /RecordDelete domain.com alpha A 192.168.1.108 /f
dnscmd DNSServer.domain.com /RecordDelete 192.168.in-addr.arpa 108.1 PTR
alpha.domain.com /f
dnscmd DNSServer.domain.com /RecordAdd domain.com alpha A 116.123.2.47
dnscmd DNSServer.domain.com /RecordAdd 116.123.in-addr.arpa 47.2 PTR alpha.domain.com
```

You would add the following to your failback script to delete the host and reverse lookup entries and add new entries associating the source with its original identity.

```
dnscmd DNSServer.domain.com /RecordDelete domain.com alpha A 116.123.2.47 /f
dnscmd DNSServer.domain.com /RecordDelete 116.123.in-addr.arpa 47.2 PTR
alpha.domain.com /f
dnscmd DNSServer.domain.com /RecordAdd domain.com alpha A 192.168.1.108
dnscmd DNSServer.domain.com /RecordAdd 192.168.in-addr.arpa 108.1 PTR alpha.domain.com
```

See your Windows documentation or the Microsoft web site for more details on the DNSCMD command.

Double-Take Availability DFO utility

DNS updates can be added to your failover and failback scripts by using the Double-Take Availability DFO utility as long as the utility has been registered and the proper privileges are configured.

1. Extract the DFO utility to the location where the Double-Take Availability program files are installed on the target.
2. From a command prompt, change to the Double-Take Availability program files directory and register the DFO utility by entering the command regsvr32 capicom.dll
3. Create a user account that has full control on the WMI DNS namespace on the source's primary DNS server.
    a. Select **Start**, **Run**, and enter the command mmc.
    b. After the Microsoft Management Console starts, select **File**, **Add/Remove Snap-in**.
    c. Click **Add**, select **WMI Control**, click **Add** again, confirm the local computer is selected, and then click **Finish**.
    d. Close the snap-in dialog box and then click **OK** to return to the console.
    e. Right-click **WMI Control** and select **Properties**.
    f. On the **Security** tab, expand the tree under **Root**.
    g. Select **MicrosoftDNS** and click **Security**.
    h. Click **Add** and identify the user account that you want the DFO utility to use.
    i. Grant the user account permissions for Execute Methods, Full Write, Partial Write, Provider Write, Enable Account, Remote Enable, and Read Security.
    j. Click **OK** to close all open dialog boxes and then close the console.
    k. Restart the Windows Management Instrumentation service for the changes to take effect. .
4. Add the same user account that has full control on the WMI DNS namespace to the domain's DnsAdmins group where the source's primary DNS server is located.
    a. Select **Start**, **Programs**, **Administrative Tools (Common)**, and **Active Directory Users and Computers**.
    b. Right-click the **DnsAdmins** group and select **Properties**.
    c. Select the **Members** tab, click **Add**, and identify the user account that you granted full control on the WMI DNS namespace.
    d. **Click** OK to close all open dialog boxes and then close Active Directory Users and Computers.
5. Add the appropriate DFO command to your failover script using the following

syntax.

Command

DFO

Description

Used in scripts to failover DNS server name

Syntax

DFO [/DNSSRVNAME <*dns_server_name*>] /SRCNAME <*source_fqd_name*> /SRCIP <*source_ip*> /TARIP <*target_ip*> /TARNAME <*target_fqd_name*> [/RECORDTYPE <*rec_type*>] [/USERNAME <*user_name*> /PASSWORD <*password*>] [/DNSZONE <*zone_name*>] [/DNSDOMAIN <*domain_name*>] [/LOGFILE <*file_name*>] /FAILBACK [*fb_switch*] [/SETPASSWORD <*user_name*> <*password*>[*machine*][*file*]] [/GETPASSWORD] [/LOCK] [/UNLOCK] /TRUSTEE <*trustee_name*> [/VERBOSE] [/FLUSHDNS /MACHINE <*machine_fqd_name*>] [/TTL <*seconds*>] [/ADDOMAIN <*active_directory_domain_name*> [/SOURCEDN <*source_domain_name*> [/TEST] [/DEBUG] [/HELP]

Options

- DNSSRVNAME *dns_server_name*—The name of the source domain/zone's primary DNS server. If not specified, the local machine will be used.
- SRCNAME *source_fqd_name*—The source machine's fully qualified domain name
- SRCIP *source_ip*—The source machine's IP address
- TARIP *target_ip*—The target machine's IP address
- TARNAME *target_fqd_name*—The target machine's fully qualified domain name (required only for failback)
- RECORDTYPE *rec_type*—The type of DNS resource records to modify or list. Values record types are ALL, MSEXCHANGE, A, CNAME, MX, PTR, STD, or STANDARD. STD and STANDARD are used to specify a non-Exchange record (minus the MX records). By default, all record types are included.
- USERNAME *user_name*—The domain name of the user account. If not specified, the account running the utility will be used.

- PASSWORD ***password***—The password associated with the user account
- DNSZONE ***zone_name***—The name of the DNS zone or DNS container, used to refine queries
- DNSDOMAIN ***domain_name***—The name of the DNS domain, used to refine queries
- LOGFILE ***file_name***—The name of the log file
- FAILBACK ***fb_switch***—Denotes a failback procedure, performed after a failed source is recovered or restored (required for failback). By default, the DFO will only failback records in the dfo_failback_config.dat file. This option allows you to enter search criteria to identify the records to change back, even if they are not in the configuration file. This option is also used if the dfo_failback_config.dat file is missing.
- SETPASSWORD ***user_name password machine file***—Stores user credentials on the specified machine in the specified file for later use. The file will be encrypted. This option must be run separately from a modify or list activity.
- GETPASSWORD—Retrieves previously stored user credentials. This option can only be used if the credentials were previously stored with the setpassword option.
- LOCK—Allows Active Directory locking for the A record type of the source specified without modifying the record
- UNLOCK—Allows Active Directory unlocking for the A record type of the source specified without modifying the record
- TRUSTEE ***trustee_name***—The domain account for the source machine (domain\machine$). DFO attempts to deny write permissions to the DNS A record on failover for the account identified as the trustee. “Deny write permissions” is then removed from the DNS A record on failback. This keeps the source server from reclaiming its DNS A record if it comes back online prior to failback.
- VERBOSE—Logging and display level set to maximum detail
- FLUSHDNS MACHINE ***machine_fqd_name***—Runs the ipconfig /flushdns command to flush the DNS cache on the specified machine. Use the fully-qualified domain name of the machine.
- TTL ***seconds***—Specifies the number of seconds for the time to live value of all modified records

- ADDOMAIN ***active_directory_domain_name***—The name of the Active Directory domain
- SOURCEDN ***source_domain_name***—The name of the source's domain
- TEST—Runs in test mode so that modifications are not made, only listed
- DEBUG—Forces DFO to write the DNS resource record as-is to the dfolog.log file prior to any DFO modify or list activity.
- HELP—Displays the syntax of the DNS Failover utility

Examples

- dfo /dnssrvname gamma.domain.com /srcname alpha.domain.com /srcip 206.31.4.10 /verbose (Lists all resource records on the specified DNS server that match the source criteria)
- dfo /dnssrvname gamma.domain.com /srcname alpha.domain.com /srcip 206.31.4.10 /tarip 210.11.12.13 /verbose (Modifies all resource records on the specified DNS server that match the source criteria, using the credentials of the account running the utility to connect to the DNS server)
- dfo /dnssrvname gamma.domain.com /srcname alpha.domain.com /srcip 206.31.4.10 /tarip 210.11.12.13 /username domain.com\admin /password /verbose (Modifies all resource records on the specified DNS server that match the source criteria, using the username and password to connect to the DNS server)
- dfo /dnssrvname gamma.domain.com /srcname alpha.domain.com /srcip 210.11.12.13 /tarname beta.domain.com /tarip 206.31.4.10 /failback /verbose (Fails back all resource records on the specified DNS server that were changed on failover)
- dfo /setpassword domain.com\admin password (Stores the user name and password in an encrypted file)
- dfo /dnssrvname gamma.domain.com /srcname alpha.domain.com /srcip 206.31.4.10 /tarip 210.11.12.13 /username domain.com\admin /getpassword /verbose (Modifies all resource records on the specified DNS server that match the source criteria, using the specified username and retrieving the password from the encrypted file)

Non-Microsoft DNS

If you are using a non-Microsoft DNS server (such as Unix) or if you are in a non-domain configuration (such as a workgroup), when Double-Take Availability failover occurs, DNS is not automatically updated. If the end-users use DNS to resolve server names and the source IP address was not failed over to the target, additional DNS updates will be required because the host records for the source will remain intact after failover. You can automate this process by scripting the DNS updates in the failover and failback scripts.

If you are protecting an application, you can configure the Application Manager not to check for Microsoft DNS or you can start the Application Manager from the command line and specify the /altdns option after your application switch.

For other workload protections, one option is to use a BIND DNS client for DNS scripting. The following steps provide an example of how you can use a BIND DNS client for DNS failover and failback scripting. You may need to modify this example to fit your environment.

1. Go to www.isc.org and download the appropriate BIND DNS client.
2. Install the BIND client on the target server.
3. Set a PATH statement for the BIND directory to ensure that it runs every time the executable is called.
4. Create a failover script file in the Double-Take Availability directory.
5. Add the following line to the failover script file, substituting your Double-Take Availability directory for install_location.

   ```
   nsupdate.exe "c:\install_location\dnsover.txt"
   ```

6. Save the failover script file.
7. Create a text file, called dnsover.txt in the Double-Take Availability directory.
8. Add the following lines to the dnsover.txt file, substituting your source name, fully-qualified domain name, target name, and target IP address as appropriate.

   ```
   update delete source_server_name.fully_qualified_domain_name.com A
   update add target_server_name.fully_qualified_domain_name.com 86400 A
   target_server_IP_address
   send
   ```

9. Save the dnsover.txt file.
10. Create a failback script file in the Double-Take Availability directory.

11. Add the following line to the failback script file, substituting your Double-Take Availability directory for install_location.

```
nsupdate.exe "c:\install_location\dnsback.txt"
```

12. Save the failback script file.
13. Create a text file, called dnsback.txt in the Double-Take Availability directory.
14. Add the following lines to the dnsback.txt file, substituting your target name, fully-qualified domain name, source name, and source IP address as appropriate.

```
update delete target_server_name.fully_qualified_domain_name.com A
update add source_server_name.fully_qualified_domain_name.com 86400 A
source_server_IP_address
send
```

15. Save the dnsback.txt file.
16. Change the Double-Take service on the target server to a domain account that has rights to modify BIND DNS. Stop and start the service to have it take effect.

Macintosh shares

A share is any volume, drive, or directory resource that is shared across a network. During failover, the target can assume or add any source shares so that they remain accessible to the end users. Automatic share failover only occurs for standard Windows file system shares. Other shares, including Macintosh volumes, must be configured for failover through the failover scripts or created manually on the target.

1. On your target, set the File Server for Macintosh service to manual startup. This allows the post-failover script on the target to control when the service starts on the target.
2. Create each volume on the target machine exactly as it exists on the source. Use the Shared Folder wizard to configure each volume as a Macintosh-accessible volume. Follow these steps to start the wizard.
    a. Open the Control Panel and click **Administrative Tools**.
    b. Select **Configure Your Server**.
    c. In the Configure Your Server window, click the **File Server** link.
    d. Click **Start the Shared Folder wizard** to start the wizard, and then follow the directions provided by the wizard. On the Create Shared Folders screen, you must enable **Apple Macintosh**.

    Note: You can automate the creation of the volumes during the failover process by using the macfile volume command in the post-failover batch file. For detailed information on how to use this command, see your Windows reference guide.

    ---

3. On the target machine, copy the chngname utility, chngname.exe, from the \tools\Win2K directory of the Double-Take Availability CD or from the Double-Take Software support web site to the directory where Double-Take Availability is installed.
4. Add the following to your failover script.

```
rem Commands for Macintosh-accessible volume failover
rem The chngname utility (chngname.exe) must be located in the same
rem directory where Double-Take Availability is installed.
rem The following command temporarily changes the name of the server.
You
rem will need to replace <drive>:\<directory>\ with the location of
rem your Double-Take Availability chngname utility and replace
rem source_name with the name of the source machine.
```

```
<drive>\<directory>\chngname /s source_name
rem The following command starts the File Server for Macintosh service
net start "File server for Macintosh"
rem The following command changes the name of the server back to its
rem original name. You will need to replace <drive>:\<directory>\ with
rem the location of your Double-Take Availability chngname utility.
<drive>\<directory>\chngname /t
```

In the event of a failure, the Macintosh clients must remap the volume in order to access it. From the Macintosh client, use the Chooser to select the volume that needs to be remapped.

NFS Shares

A share is any volume, drive, or directory resource that is shared across a network. During failover, the target can assume or add any source shares so that they remain accessible to the end users. Automatic share failover only occurs for standard Windows file system shares. Other shares, including NFS shares, must be configured for failover through the failover scripts or created manually on the target.

1. On your target, set the NFS service to manual startup. This allows the post-failover script on the target to control when the service starts on the target.
2. Create each shared drive or directory on the target exactly as it exists on the source. Configure each drive or directory as an NFS share by following these steps.
    a. Right-click the drive or directory that you want to share and select **Sharing**.
    b. Click the **NFS Sharing** tab on the Program Files Properties dialog box.
    c. Enable **Share this folder**, provide the name of the share, and click **OK** to share the folder as an NFS share.
3. On the target machine, copy the chngname utility, chngname.exe, from the \tools\Win2K directory of the Double-Take Availability CD or from the Double-Take Software support web site to the directory where Double-Take Availability is installed.
4. Add the following to your failover script.

```
rem Commands for NFS share failover
rem The chngname utility (chngname.exe) must be located in the same
rem directory where Double-Take Availability is installed.
rem The following command temporarily changes the name of the server.
You
rem will need to replace <drive>:\<directory>\ with the location of
rem your Double-Take Availability chngname utility and replace
rem source_name with the name of the source machine.
<drive>\<directory>\chngname /s source_name
rem The following command starts the NFS service
net start "Server for NFS"
```

In the event of a failure, the clients must remount the shares in order to access them.

Workload monitoring

Once a workload protection is established you will want to monitor the protection. You can monitor the workload protection using the same Double-Take Availability client that you used to establish the workload protection, or you can use several general monitoring tools that are available for all workload types.

- Data workloads
- Full-server workloads
- Application workload
- Virtual workloads
- Cluster workloads
- Log files
- Windows Event messages
- Statistics
- Performance Monitor
- SNMP
- Error codes

Data workloads

When you are working with data workloads, you can monitor the connection and you can monitor the status of failover monitoring.

- Monitoring a data workload
- Monitoring failover monitoring

Monitoring a data workload

When a source is highlighted in the left pane of the Replication Console, the connections and their statistics are displayed in the right pane. Additionally, colors and icons are used for the connections, and the Double-Take Availability servers, to help you monitor your connections.

- Connection statistics
- Connection and server display

Connection statistics

1. You can change the statistics that are displayed by selecting **File**, **Options** and selecting the **Statistics** tab.
2. The statistics displayed in the Replication Console will be listed with check boxes to the left of each item. Mark the check box to the left of each statistic that you want to appear, and clear the check box to the left of each statistic that you do not want to appear.
3. The statistics appear on the Replication Console in the order they appear on the **Statistics** tab. If you want to reorder the statistics, highlight the statistic to be moved and select the up or down arrow button, to the right of the vertical scroll bar, to move the selection up or down in the list. Repeat this process for each statistic that needs to be moved until you reach the desired order.
4. If you have made changes to the statistics list and have not yet saved them, you can go back to the previously used settings by clicking **Reset to Last**. This will revert the list back to the last saved settings.
5. To return the statistics list to the Double-Take Availability default selection and order, click **Reset to Default**.
6. Click **OK** to apply and save any changes that have been made to the order or display of the Replication Console statistics.

Statistics marked with an asterisk (*) are not displayed, by default.

Replication Set

Replication set indicates the name of the connected replication set.

Connection ID

The connection ID is the incremental counter used to number each connection established. This number is reset to one each time the Double-Take service is restarted.

Target Name

The name of the target as it appears in the server tree in the left pane of the Replication Console. If the server's name is not in the server tree, the IP address will be displayed.

Target IP

The target IP is the IP address on the target machine where the mirroring and replication data is being transmitted.

Target Data State

- **OK**—The data on the target is in a good state.
- **Mirroring**—The target is in the middle of a mirror process. The data will not be in a good state until the mirror is complete.
- **Mirror Required**—The data on the target is not in a good state because a remirror is required. This may be caused by an incomplete or stopped mirror or an operation may have been dropped on the target.
- **Restore required**—The data on the source and target do not match because of a failover condition. Restore the data from the target back to the source. If you want to discard the changes on the target, you can remirror to resynchronize the source and target.
- **Snapshot reverted**—The data on the source and target do not match because a snapshot has been applied on the target. Restore the data from the target back to the source. If you want to discard the changes on the target, you can remirror to resynchronize the source and target.

Target Status

- **OK**—The target machine is active and online.
- **Not Loaded**—The target module is not loaded on the target. (For example, the activation code is invalid.)
- **Paused**—The target machine is paused by user intervention.
- **Retrying**—The target machine is retrying operations for the connection.

This field may not be updated until there is source/target activity.

Commit Mode *

The commit mode status indicates the connection status.

- **Real-time**—Data is being transmitted to the target machine in real-time.
- **Scheduled**—Data is waiting to be transmitted to the target machine until one or more transmit options have been met.

Transmit Mode

- **Started**—Data is being transferred to the target machine.
- **Paused**—If the transmission is real-time and the transmission has been paused, the **Transmit Mode** indicates **Paused**.
- **Scheduled**—If the transmission is scheduled, the **Transmit Mode** indicates **Scheduled**.
- **Stopped**—Data is not being transferred to the target machine.
- **Error**—There is a transmission error.

Mirror Status

- **Mirroring**—If the file size of the replication set has not been calculated and the data is being mirrored to the target machine, the **Mirror Status** will indicate **Mirroring**.
- **Idle**—Data is not being mirrored to the target machine.
- **Paused**—Mirroring has been paused.
- **Percentage Complete**—If the file size of the replication set has been calculated and the data is being mirrored to the target machine, the **Mirror Status** will display the percentage of the replication set that has been sent.
- **Waiting**—Mirroring is complete, but data is still being written to the target.
- **Restoring**—Data is being restored from the target to the source.
- **Verifying**—Data is being verified.
- **Removing Orphans**—Double-Take Availability is checking for orphan files within the target path location (files that exist on the target but not on the source). These files will be removed.
- **Archiving**—Data is being archived or an archive report is being run.

Replication Status

- **Replicating**—Data is being replicated to the target machine.
- **Ready**—There is no data to replicate to the target machine.
- **Stopped**—Replication has stopped.
- **Pending**—If auto-remirror is enabled and you have experienced a source or target failure and recovery, the status will change to pending while the connections are reestablished and will update when the remirror begins. If auto-remirror is disabled and you have experienced a source or target failure and recovery, replication will be Pending until a remirror is performed. Without a remirror, data integrity cannot be guaranteed.
- **Out of Memory**—Kernel memory has been exhausted.
- **Failed**—TheDouble-Take service is not receiving replication operations from the Double-Take Availability driver. Check the Double-Take Availability log and Event Viewer for driver related issues.

Queued (Ops) *

The queued (ops) statistic indicates the total number of mirror and replication operations that are in the source queue.

Sent (Bytes)

The sent (bytes) statistic indicates the total number of mirror and replication bytes that have been transmitted to the target.

Sent Compressed (Bytes)

The sent compressed (bytes) statistic indicates the total number of compressed mirror and replication bytes that have been transmitted to the target. If compression is disabled, this statistic will be the same as sent (bytes).

Intermediate Queue (Bytes) *

The intermediate queue (bytes) indicates the total amount of memory being used by the operations buffer queue.

Disk Queue (Bytes)

The disk queue (bytes) indicates the amount of disk being used to queue data on the source.

Queued Replication (Bytes)

The queued replication (bytes) statistic is the total number of replication bytes that are remaining to be transmitted from the source.

Sent Replication (Bytes)

The sent replication (bytes) statistic is the total number of replication bytes that have been transmitted to the target.

Sent Compressed Replication (Bytes) *

The sent compressed replication (bytes) statistic is the total number of compressed replication bytes that have been transmitted to the target. If compression is disabled, this statistic will be the same as sent replication (bytes).

Queued Mirror (Ops) *

The queue mirror (ops) statistic is the total number of mirror operations in the queue.

Sent Mirror (Bytes)

The sent mirror (bytes) statistic is the total number of mirror bytes that have been transmitted to the target.

Sent Compressed Mirror (Bytes) *

The sent compressed mirror (bytes) statistic is the total number of compressed mirror bytes that have been transmitted to the target. If

compression is disabled, this statistic will be the same as sent mirror (bytes).

Skipped Mirror (Bytes)

The skipped mirror (bytes) statistic is the total number of bytes that have been skipped when performing a difference or checksum mirror. These bytes are skipped because the data is not different on the source and target machines.

Remaining Mirror (Bytes)

The remaining mirror (bytes) statistic is the total number of mirror bytes that are remaining to be sent to the target.

Queued Replication (Ops) *

The queued replication (ops) statistic is the total number of replication operations in the queue.

Last File Touched

The last file touched identifies the last file that Double-Take Availability transmitted to the target. If you are using long file names (more than several thousand characters long) you may want to disable the display of this statistic to improve Replication Console response times.

Connected Since

Connected since is the date and time indicating when the current connection was made. This field is blank, indicating that a TCP/IP socket is not present, when the connection is waiting on transmit options or if the transmission has been stopped. This field will maintain the date and time, indicating that a TCP/IP socket is present, when transmission has been paused.

Bandwidth Limit (Kbps)

If bandwidth limiting has been set, this statistic identifies the limit. The keyword **Unlimited** means there is no bandwidth limit set for the connection.

Connection and sever display

You can configure when the icons and colors change to accommodate your network environment. For example, a slow or busy network may need longer delays before updating the icons or colors.

1. Select **File**, **Options**. On the **Configuration** tab, you will see **Site Monitor** and **Connection Monitor**. The **Site Monitor** fields control the icons on the left pane of the Replication Console and the icons on the right pane when a group is highlighted in the left pane. The **Connection Monitor** field controls the display when a server is highlighted in the left pane. These two separate monitoring capabilities allow for flexible monitoring.
2. Under **Site Monitor**, specify **Check Status Interval** to identify the number of seconds between requests sent from the Replication Console to the servers in order to update the display. Valid values are between 0 and 3600. The default setting is 30 seconds.
3. Under **Site Monitor**, specify **Missed Status Responses** to identify the number of responses from a server that can be missed before the Replication Console considers communications lost and updates the icons. Valid values are between 1 and 100. The default setting is 2.
4. Under **Connection Monitor**, specify **Missed Status Responses** to identify the number of responses from a server that can be missed before the Replication Console considers communications lost and updates the icons and colors. Valid values are between 0 and 1000. The default setting is 5.
5. Click **OK** to save the settings.

Note: If the **Site Monitor** and **Connection Monitor** settings are different, at times, the icons and color may not be synchronized between the left and right panes.

The following icons are displayed in the left pane.

 —An icon with yellow and blue servers indicates a server that is working properly.

 —A hammer over a server indicates an activation code violation. Check the Double-Take Availability log or Event messages for more information.

 —A red X on a server icon indicates the Replication Console cannot communicate with that server or that is a problem with one of the server's connections. If the connection background is gray, it is a communication issue. If the connection also has a red X, it is a connection issue.

 —Two red vertical lines on a server icon indicates the target is paused.

 —A red tree view (folder structure) on a server icon indicates a restore is required because of a failover.

 —A black X on a server icon indicates the server is not running Double-Take Availability.

 —A yellow folder with a blue server indicates a group folder that is working properly.

 —A black exclamation point inside a yellow triangle on a group folder indicates there is a communication error on one of the servers in the group. Drill down through the group until you find the server that has the communication error.

 —A white X inside a red circle on a group folder indicates there is a connection error on one of the servers in the group. Drill down through the group until you find the server that has the connection error.

The following icons and colors are displayed in the right pane when a server is highlighted in the left pane.

 —A green checkmark on a connection indicates the connection is working properly.

 —A red X on a connection indicates a connection error. For example, an error may be caused by broken transmission or pending replication. To determine the exact problem, locate the connection data item that appears in red.

—A lock icon on a connection indicates target path blocking is enabled for the connection. This prohibits writing to the path on the target where the replication data is stored.

White background—If the connection background is white, the Replication Console and the source are communicating.

Gray background—If the connection background is gray, the Replication Console and the source are no longer communicating. The connection data stops updating

once communications have stopped. Once communications have been reestablished, the connection background will change back to white.

The following icons are displayed in the right pane when a group is highlighted in the left pane.

 —An icon with a network cable between two servers on the right pane indicates there are no established Double-Take Availability connections to this server. This icon also indicates that communications between the Replication Console and the server are working properly.

 —An icon with two servers on the right pane indicates this server has active connections that are working properly.

 —A yellow server icon with a red X on the right pane indicates a connection error. For example, an error may be caused by broken transmission or pending replication. To determine the exact problem, locate the connection data item that appears in red.

 —An icon with a network cable between two servers and marked with a red X on the right pane indicates a communication error between the Replication Console and the server.

 —A blue server icon with a red X indicates a restore is required because of a failover.

 —A network cable with a black X indicates the server is not running Double-Take Availability.

Monitoring failover monitoring

Since it can be essential to quickly know the status of failover monitoring, Double-Take Availability offers various methods for monitoring failover monitoring. When the Failover Control Center is running, you will see four visual indicators:

- The Failover Control Center Time to Fail counter
- The Failover Control Center status bar located at the bottom of the window
- The Failover Control Center colored bullets to the left of each IP address and source machine
- The Windows desktop icon tray containing a failover icon

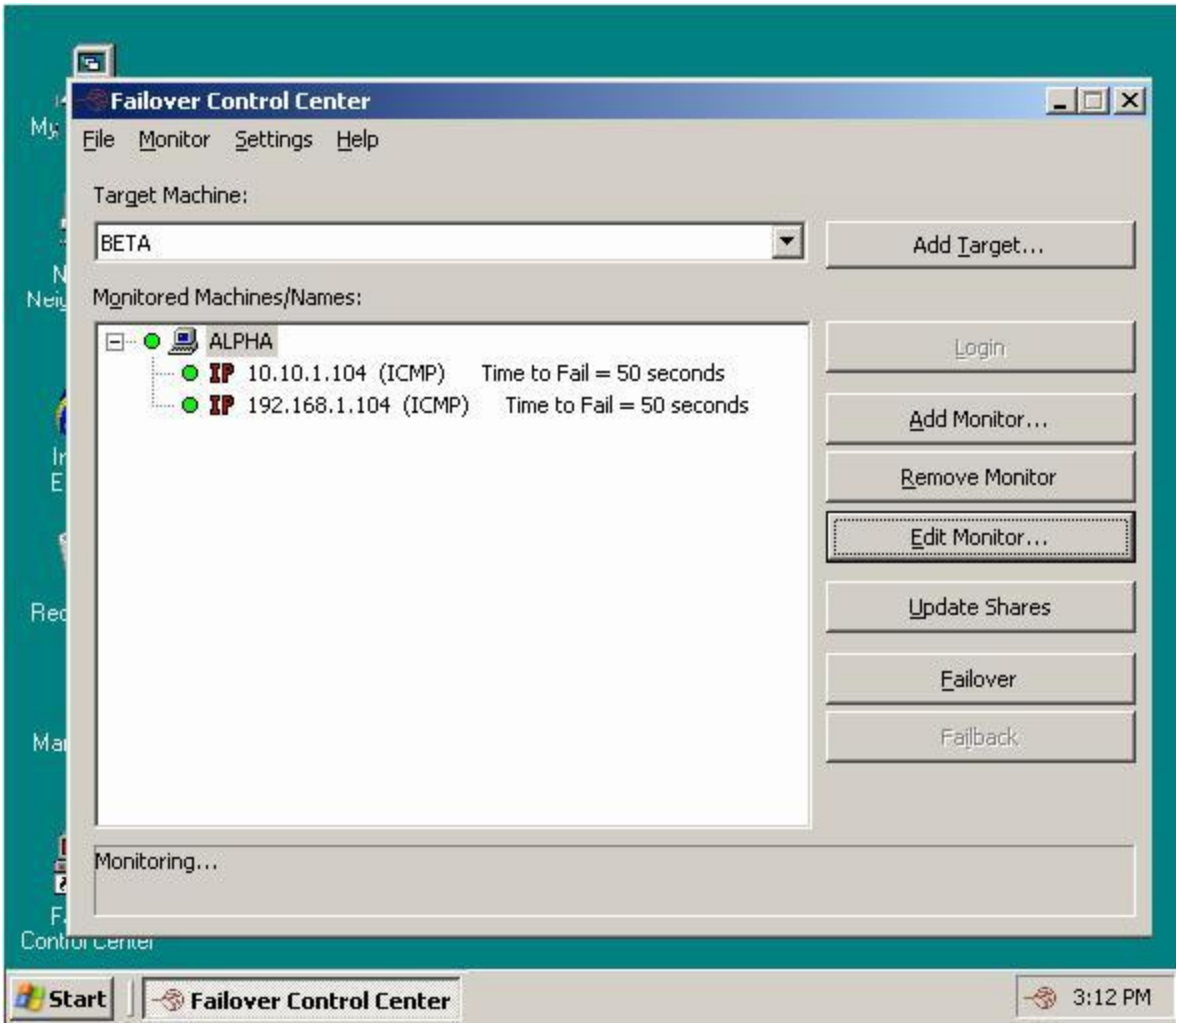

Note: You can minimize the Failover Control Center and, although it will not appear in your Windows taskbar, it will still be active and the failover icon will still appear in the desktop icon tray.

The Failover Control Center does not have to be running for failover to occur.

The following table identifies how the visual indicators change when the source is online.

Time to Fail Countdown

The Time to Fail counter is counting down and resetting each time a response is received from the source machine.

Status Bar

The status bar indicates that the target machine is monitoring the source machine.

Colored Bullets

The bullets are green.

When the Time to Fail value has decreased by 25% of the entire timeout period, the bullet changes from green to yellow, indicating that the target has not received a response from the source. The yellow bullet is a caution signal. If a response from the source is received, the countdown resets and the bullets change back to green. If the countdown reaches zero without the target receiving a response from the source, failover begins.

Desktop Icon Tray

The Windows desktop icon tray contains a failover icon with red and green computers.

The following table identifies how the visual indicators change when the source fails and failover is initiated.

Time to Fail Countdown

The Time to Fail countdown value is 0.

Status Bar

The status bar displays the source machine and IP address currently being assumed by the target.

Colored Bullets

The bullets are red.

Desktop Icon Tray

The Windows desktop icon tray contains a failover icon with red and green computers.

The following table identifies how the visual indicators change when failover is complete.

Time to Fail Countdown

The Time to Fail counter is replaced with a failed message.

Status Bar

The status bar indicates that monitoring has continued.

Colored Bullets

The bullets are red.

Desktop Icon Tray

The Windows desktop icon tray contains a failover icon with a red computer.

Monitoring a full-server workload

After you have enabled full-server protection, you can monitor the protection from the Full-Server Failover Manager, and you can review the log file generated by Full-Server Failover Manager.

The **Protection Status** is displayed in the right center of the Full-Server Failover Manager. You can tell the status of your protection from this field.

- **Disabled**—Protection for the source has not been started. The target must be validated as compatible before you can enable protection.
- **Initializing**—Double-Take Availability is initializing protection. Once initialization is complete, mirroring will automatically begin.
- **Mirroring (% complete)**—Double-Take Availability is mirroring the source's data and system state to the target. The percentage indicates how much of the mirror has been completed. Protection is not complete until the mirror has completed.
- **Mirroring (Retrying)**—You may see this status message if the target is out of disk space or if Double-Take Availability cannot write to a file on the target. Check the log file on the target for more information.
- **Mirroring (Mirror stopped)**—You may see this status message if your mirroring process has stopped. Depending on the reason for the stopped mirror, it may restart automatically. Check your Double-Take Availability log file.
- **Mirroring (Source Unavailable)**—You may see this status message if the source has become unavailable during your mirroring process. You cannot failover until the mirroring process is complete. Correct the issues causing the unavailability, and mirroring will restart automatically.
- **Mirroring (Op Dropped)**—You may see this status message if the target server cannot apply data to disk. For example, a file may be in use on the target. Check the Double-Take Availability log file on the target for more details.
- **Enabled**—The mirroring is complete and protection of the source is enabled. In the event the source should fail, the target will be able to stand-in for it.
- **Enabled (Source Unavailable)**—You may see this message when the target has lost communication with the source. If communication is reestablished before the failover monitoring time expires, the status will update to **Enabled**. If communication is not reestablished before the failover monitoring time expires, the status will update to **Failover condition met**.
- **Failover condition met**—The target has missed too many responses from the source, indicating that the source has failed. At this time, you need to manually

determine the status of the source. If the source is still up and users are accessing it, you need to resolve the communications errors between the source and target. Once the communication issue is resolved, the status will update to the appropriate state. If the source is indeed down and users are unable to access it, start failover.

- **Failing over (% complete)**—The target is in the process of failing over for the source. The percentage indicates how much of the failover has been completed.
- **Failed over**—Failover is complete. The target will automatically reboot.

In addition to the status displayed in Full-Server Failover Manager, a log file is generated detailing processing information. By default, Full-Server Failover Manager logs basic processing information. To view the log file, select **File**, **Logs**, **View Full-Server Failover Manager Log**. The log file will be opened automatically in Notepad. The log file, FFMLog.log, is located on the in the directory where you installed it. You can clear the log file by selecting **File**, **Logs**, **Clear Full-Server Failover Manager Log**.

Monitoring an application workload

After you have enabled application protection, you can monitor the protection from the Application Manager **Monitor** tab.

Protection Status

- **Unprotected**—No connection exists
- **Warning**—A connection exists but has issues
- **Protected**—A connection exists and is active
- **Synchronizing**—Mirroring is in progress
- **Unknown**—The protection status could not be determined
- **Failing over**—Failover from the source to the target is in progress
- **Failed over**—Failover is complete and the target has assumed the source role
- **Failing back**—Failback from the target to the original source is in progress
- **Restoring**—Mirroring (target to source) is in progress

Monitoring Status

- **Disabled**—Monitoring is disabled
- **Enabled**—Monitoring is active
- **Failover condition met**—The source server is unavailable
- **Failing over**—Failover from the source to the target is in progress
- **Failed over**—Failover is complete and the target has assumed the source role
- **Failing back**—Failback from the target to the original source is in progress

Mirror Status

- **Calculating**—The amount of data to be mirrored is being calculated
- **Idle**—Data is not being mirrored to the target machine
- **Mirroring**—Data is being mirrored to the target machine
- **Paused**—Mirroring has been paused
- **Removing Orphans**—Double-Take Availability is checking for orphan files within the target path location (files that exist on the target but not on the source). These files will be removed.
- **Verifying**—Data is being verified

- **Restoring**—Data is being restored from the target to the source
- **Unknown**—The mirror status could not be determined

Mirror Remaining

The percentage of the mirror remaining

Replication Status

- **Replicating**—Data is being replicated to the target machine
- **Ready**—There is no data to replicate to the target machine
- **Stopped**—Replication has stopped
- **Out of Memory**—Kernel memory has been exhausted

Transmit Mode

- **Started**—Data is being transferred to the target machine
- **Paused**—Data transmission has been paused
- **Stopped**—Data is not being transferred to the target machine.
- **Error**—There is a transmission error.

Target State

- **Online**—The target is active and online
- **Mirroring**—The target is in the middle of a mirror process. The data will not be in a good state until the mirror is complete.
- **Mirroring stopped**—The mirroring process has been stopped
- **Remirror required**—The data on the target is not in a good state because a remirror is required. This may be caused by an incomplete or stopped mirror or an operation may have been dropped on the target.
- **Retrying**—The target machine is retrying operations
- **Paused**—The target machine has been paused
- **Pausing**—The connection is pausing
- **Restore required**—The data on the source and target do not match because of a failover condition. Restore the data from the target back to the source. If you want to discard the changes on the target, you can remirror to resynchronize the source and target.
- **Replicating**—Data is being replicated to the target
- **Snapshot reverted**—The data on the source and target do not match because a snapshot has been applied on the target. Restore the data from the target back to the source. If you want to discard the changes on the target, you can remirror to resynchronize the source and target.
- **Target path blocked**—Target path blocking is enabled

- **Target path unblocked**—Target path blocking is disabled
- **Unknown**—The state could not be determined

Connected Since

The date and time indicating when the current connection was made

Transmitted

The amount of data transmitted from the source to the target

Compressed

The amount of compressed data transmitted from the source to the target

Source Queue

The amount of data in queue on the source

Target Queue

The amount of data in queue on the target

Monitoring virtual workloads

When you are working with virtual workloads, you can monitor the connection and you can monitor the status of failover monitoring. If you are protecting host-level virtual disk files (the .vmdk files) from an ESX source to an ESX target, you will need to use the Double-Take Availability for VMware Infrastructure console to monitor your protection. For all other virtual workloads, use the Double-Take Console to monitor your protection.

Monitoring virtual workloads in the Double-Take Console

Click **Monitor Connections** from the main Double-Take Console toolbar. The **Monitor Connections** page allows you to view information about your connections. You can also manage your connections from this page.

- Overview connection information displayed in the top pane
- Filtering the connections displayed in the top pane
- Detailed connection information displayed in the bottom pane
- Connection controls available in the bottom pane

Overview connection information displayed in the top pane

The top pane displays high-level overview information about your connections.

Column 1 (Blank)

The first blank column indicates the state of the connection.

 The connection is in a healthy state.

 The connection is in a warning state.

 The connection is in an error state.

 The connection is in an unknown state.

Connection

The name of the connection

Source Server

The name of the source

Target Server

The name of the target

Replication Set

The name of the replication set

Activity

There are many different **Activity** messages that keep you informed of the connection activity. Most of the activity messages are informational and do not require any administrator interaction. If you see error messages, check the connection details.

Mirror Status

- **Calculating**—The amount of data to be mirrored is being calculated.
- **Mirroring**—If the file size of the replication set has not been calculated and the data is being mirrored to the target machine, the **Mirror Status** will indicate Mirroring.
- **Percentage Complete**—If the file size of the replication set has been calculated and the data is being mirrored to the target machine, the **Mirror Status** will display the percentage of the replication set that has been sent.

- **Waiting**—Mirroring is complete, but data is still being written to the target.
- **Idle**—Data is not being mirrored.
- **Paused**—Mirroring has been paused.
- **Stopped**—Mirroring has been stopped.
- **Removing Orphans**—Orphan files on the target are being removed or deleted depending on the configuration.
- **Verifying**—Data is being verified between the source and target.
- **Restoring**—Data is being restored from the target to the source.
- **Archiving**—Data is being archived or an archive report is being run.
- **Unknown**—The console cannot determine the status.

Replication Status

- **Replicating**—Data is being replicated to the target.
- **Ready**—There is no data to replicate.
- **Pending**—Replication is pending.
- **Stopped**—Replication has been stopped.
- **Out of Memory**—Replication memory has been exhausted.
- **Failed**—The Double-Take service is not receiving replication operations from the Double-Take driver. Check the Event Viewer for driver related issues.
- **Unknown**—The console cannot determine the status.

Transmit Mode

- **Started**—Data is being transmitted to the target.
- **Paused**—Data transmission has been paused.
- **Scheduled**—Data transmission is waiting on schedule criteria.
- **Stopped**—Data is not being transmitted to the target.
- **Error**—There is a transmission error.
- **Unknown**—The console cannot determine the status.

Filtering the connections displayed in the top pane

You can filter the connections displayed in the top pane using the toolbar buttons in that pane.

View Connection Details 

Leave the **Monitor Connections** page and open the View Connection Details page

Filter

Select a filter option from the drop-down list to only display certain connections. You can display **Healthy Connections**, **Connections with warnings**, or **Connections with errors**. To clear the filter, select **All connections**.

View Connections with Warnings 

Display only connections with warnings

View Connections with Errors 

Display only connections with errors

Type a server name

Only those servers that contain the entered text will be displayed

---

Detailed connection information displayed in the bottom pane

The details displayed in the bottom pane of the **Monitor Connections** page will depend on the type of protection workload that is highlighted in the top pane.

Name

The name of the server

Activity

There are many different **Activity** messages that keep you informed of the connection activity. Most of the activity messages are informational and do not require any administrator interaction. If you see error messages, check the connection details.

Source server

The name of the source

Target server

The name of the target

Bytes sent

The total number of mirror and replication bytes that have been transmitted to the target

Bytes sent compressed

The total number of compressed mirror and replication bytes that have been transmitted to the target. If compression is disabled, this statistic will be the same as **Bytes sent**.

Connected since

The date and time indicating when the current connection was made. This field is blank, indicating that a TCP/IP socket is not present, when the connection is waiting on transmit options or if the transmission has been stopped. This field will maintain the date and time, indicating that a TCP/IP socket is present, when transmission has been paused.

Protection type

The type of workload protection

Hypervisor

The type of target virtual server host

Protection status

The status of the workload protection

Protected volumes

The volumes that are being protected

Target datastore or **Target path**

The location on the target where the source replica is being stored

Automatic failover

Indicates if failover will be automatic and the number of retries that have been attempted if the source is unresponsive

---

Connection controls available in the bottom pane

The connection controls available in the bottom pane of the **Monitor Connections** page will depend on the type of protection workload that is highlighted in the top pane.

Configure 

Opens the Protection Summary page

Delete 

Removes configuration information for the selected protection

If you no longer want to protect the source and no longer need the replica of the source on the target, select the appropriate delete option when prompted. The option name will vary depending on your workload type. Selecting this option will remove the connection and completely delete the replica virtual machine on the target.

If you no longer want to mirror and replicate data from the source to the target but still want to keep the replica of the source on the target, select the appropriate keep option when prompted. The option name will vary depending on the your workload type. For example, you may want to use this option to relocate the virtual hard disks and create a new job between the original source and the new location. Selecting this option, will preserve and register the source replica on the target, provided it has been fully synchronized. If the source replica is not fully synchronized, related files will be kept on the target but will not be registered.

Start protection 

Enables protection for the selected protection

If you have previously stopped protection, the virtual hard disks on the target will be checked. If they are the same as the source, replication only (no mirroring) will begin. If they are not the same but there is a file on the target, a difference mirror will begin.

If you have previously paused protection, the protection job will resume where it left off.

Pause protection 

Pause the selected protection

Stop protection 

Stops the selected protection

Failover 

Initiate failover by shutting down the source and starting the replica of the source on the target

Undo Failover 

For some workload types, you can undo failover after it has occurred. This resets the servers and the job back to their original state.

Reverse protection 

For some workload types, you can reverse the protection. The connection will start mirroring in the reverse direction with the connection name and log file names changing accordingly. After the mirror is complete, the job will continue running in the opposite direction.

View protection error 

Displays the most recent error associated with the selected protection

Viewing connection details

The **View Connection Details** page allows you to view information about a specific connection.

Connection name

The name of the connection

Description

The connection status

Health

The connection is in a healthy state.

The connection is in a warning state.

The connection is in an error state.

The connection is in an unknown state.

Activity

There are many different **Activity** messages that keep you informed of the connection activity. Most of the activity messages are informational and do not require any administrator interaction. If you see error messages, check the rest of the connection details.

Replication set

The name of the replication set

Connection ID

The incremental counter used to number each connection established. This number is reset to one each time the Double-Take service is restarted.

Transmit mode

- **Started**—Data is being transmitted to the target.
- **Paused**—Data transmission has been paused.
- **Scheduled**—Data transmission is waiting on schedule criteria.
- **Stopped**—Data is not being transmitted to the target.
- **Error**—There is a transmission error.
- **Unknown**—The console cannot determine the status.

Target data state

- **OK**—The data on the target is in a good state.
- **Mirroring**—The target is in the middle of a mirror process. The data will not be in a good state until the mirror is complete.
- **Mirror Required**—The data on the target is not in a good state because a remirror is required. This may be caused by an incomplete or stopped mirror or an operation may have been dropped on the target.
- **Restore required**—The data on the source and target do not match because of a failover condition. Restore the data from the target back to the source. If you want to discard the changes on the target, you can remirror to resynchronize the source and target.
- **Snapshot reverted**—The data on the source and target do not match because a snapshot has been applied on the target. Restore the data from the target back to the source. If you want to discard the changes on the target, you can remirror to resynchronize the source and target.

Target route

The IP address on the target used for Double-Take Availability transmissions.

Compression

- **On** / **Level**—Data is compressed at the level specified
- **Off**—Data is not compressed

Bandwidth limit

If bandwidth limiting has been set, this statistic identifies the limit. The keyword **Unlimited** means there is no bandwidth limit set for the connection.

Connected since

The date and time indicating when the current connection was made. This field is blank, indicating that a TCP/IP socket is not present, when the connection is waiting on transmit options or if the transmission has been stopped. This field will maintain the date and time, indicating that a TCP/IP socket is present, when transmission has been paused.

Mirror status

- **Calculating**—The amount of data to be mirrored is being calculated.

- **Mirroring**—If the file size of the replication set has not been calculated and the data is being mirrored to the target machine, the **Mirror Status** will indicate Mirroring.
- **Percentage Complete**—If the file size of the replication set has been calculated and the data is being mirrored to the target machine, the **Mirror Status** will display the percentage of the replication set that has been sent.
- **Waiting**—Mirroring is complete, but data is still being written to the target.
- **Idle**—Data is not being mirrored.
- **Paused**—Mirroring has been paused.
- **Stopped**—Mirroring has been stopped.
- **Removing Orphans**—Orphan files on the target are being removed or deleted depending on the configuration.
- **Verifying**—Data is being verified between the source and target.
- **Restoring**—Data is being restored from the target to the source.
- **Archiving**—Data is being archived or an archive report is being run.
- **Unknown**—The console cannot determine the status.

Mirror percent complete

The percentage of the mirror that has been completed

Mirror remaining

The total number of mirror bytes that are remaining to be sent from the source to the target

Mirror skipped

The total number of bytes that have been skipped when performing a difference or checksum mirror. These bytes are skipped because the data is not different on the source and target.

Replication status

- **Replicating**—Data is being replicated to the target.
- **Ready**—There is no data to replicate.
- **Pending**—Replication is pending.
- **Stopped**—Replication has been stopped.
- **Out of Memory**—Replication memory has been exhausted.

- **Failed**—The Double-Take service is not receiving replication operations from the Double-Take driver. Check the Event Viewer for driver related issues.
- **Unknown**—The console cannot determine the status.

Replication queue

The total number of replication bytes in the source queue

Disk queue

The amount of disk space being used to queue data on the source

Bytes sent

The total number of mirror and replication bytes that have been transmitted to the target

Bytes sent compressed

The total number of compressed mirror and replication bytes that have been transmitted to the target. If compression is disabled, this statistic will be the same as **Bytes sent**.

Monitoring virtual workloads in the Double-Take Availability for VMware Infrastructure console

Select **Monitor protection** from the left pane of the console. The **Monitor protection** page allows you to view information about your connections. You can also manage your connections from this page.

- Overview connection information displayed in the top pane
- Detailed connection information displayed in the bottom pane
- Connection controls

Overview connection information displayed in the top pane

The top pane displays high-level overview information about your connections.

Name

The name of the connection

Status

A description of the current status of the protection

Bytes Pending

The remaining amount of data (.vmdk files plus snapshot files) that needs to be transmitted

Remaining Interval

The amount of time until the next replication cycle

Detailed connection information displayed in the bottom pane

When the **View protection details** button in the toolbar is toggled on, the bottom pane displays detailed connection information.

Source, Virtual Disk

A list of the virtual disks being protected

Source, Snapshot Data Size

The size of the snapshot of each virtual disk

Source, Last Modified

The last time the snapshot was updated

Source, Virtual machine name

The name of the virtual machine that contains the virtual disk being protected

Source, Snapshot datastore

The name of the datastore on the source storing the snapshot data

Source, Free space

The amount of free space on the source snapshot datastore

Target, Virtual machine name

The name of the target replica virtual machine

Target, Datastore

The name of the target datastore

Target, Free space

The amount of free space on the target datastore

Target, Last replication

The last time snapshot files were replicated to the target

Target, Last synchronization

The last time .vmdk files were mirrored to the target

Connection controls

The connection controls are available in the toolbar of the **Monitor protection** page.

 —Opens the **Protection Summary** page allowing you to modify some protection settings.

 —Deletes the selected protection. You will be prompted to keep or delete the associated replica virtual machine on the target. If you do not need the replica on the target, you can delete it. However, if you want to keep the replica on the target, for example, if you want to being using the replica on the target as the production server, you can keep the replica. In this case, the replica will be preserved and registered (as long as the initial mirror has completed) allowing the virtual machine to be available in VirtualCenter. If the initial mirror has not completed, the files will be available on the target ESX server but will not be registered.

 —Starts protection

 —Stops protection

 —Initiates failover by stopping the virtual machine on the source and starting the replica virtual machine on the target

 —Initiate reverse protection by mirroring and replicating from the replica virtual machine on the target to the virtual machine on the source

 —Initiate undo failover to reset the virtual machines and the protection job back to the original state

 —View errors for the selected protection

 —View details for the selected protection

---

Monitoring a cluster workload

In a standard cluster configuration, where a single copy of data resides on a SCSI disk that is shared between cluster nodes, the Double-Take Source Connection resource keeps the data synchronized between your source and target. Use the standard Windows cluster tools to monitor the status of the resource.

In a GeoCluster configuration, where data is stored on volumes local to each node and replicated to each node in the cluster, the GeoCluster Replicated Disk resource keeps the data synchronized between nodes of the cluster. You should also use the standard Windows cluster tools to monitor the status of the resource, but you should also use the following information to help you monitor the GeoCluster Replication Disk resource.

Resolving an online pending GeoCluster Replicated Disk resource

When the GeoCluster Replicated Disk resource is in an online pending state, you are protected from possible data corruption. If you are using Windows 2003, review the description of the GeoCluster Replicated Disk Status resource to see why the GeoCluster Replicated Disk resource is in the online pending state. If you are using Windows 2008, you can see the online pending status directly in the description of the GeoCluster Replicated Disk resource. If the pending state were bypassed, the node where you are trying to bring the resource online would have incomplete data, which would then be replicated to the other nodes in the cluster. This state safeguards you from corrupting your data.

There are different options for resolving an online pending state, depending on whether your operating system supports snapshots. Therefore, some of the following options may not be displayed or may be disabled if they are not valid for your configuration.

If you are using Windows 2003, right-click on the online pending resource and select the desired control. The controls are described in the following tables.

If you are using Windows 2008, right-click the online pending resource, select **Properties**, select the **Online Pending** tab, and click the desired control. The controls are described in the following tables.

Windows 2003 Menu

Revert to snapshot

Windows 2008 Menu

Revert snapshot

Description

If you have a snapshot of the target data available, you can revert to that data. If you revert to a snapshot, any data changes made after the snapshot's specified date and time will be lost. A Double-Take Availability connection will be established to replicate the node's data (at the snapshot point-in-time) to the other nodes.

Windows 2003 Menu

Discard target queue

Windows 2008 Menu

Flush Target

Description

If you have data in the target queue, you can discard that data. If you discard the queued data, you will lose the changes associated with that data made on the previously owning node. A Double-Take Availability connection will be established to replicate the node's data (without the data that was in queue) to the other nodes.

Windows 2003 Menu

Force Resource Offline

Windows 2008 Menu

Fail Resource

Description

If you are using Windows 2003, you can force the resource offline. If you are using Windows 2008, you can fail the resource. In either case, no Double-Take Availabilityconnection will be established.

Windows 2003 Menu

Verify Group

Windows 2008 Menu

Test Data

Description

With this option and snapshot capability, you can test the data on the node before deciding whether to use it. If you select this option, a snapshot of the node's current Double-Take Availability data will be taken, but the GeoCluster Replicated Disk resource does not come online, allowing you to check the data. (This means there is no Double-Take Availability connection established at this time.) Once the snapshot is taken, you can test the data on the node to see if it is viable. Once you have tested the data, you need to right-click on the online pending resource again and accept or reject the data.

Windows 2003 Menu

Accept Data

Windows 2008 Menu

Accept

Description

If you accept the data, the current data on the node will be used, and a Double-Take Availability connection will be established to replicate the current node's data to the other nodes. If any other nodes in the cluster contain more recent data, this node will overwrite that data and it will be lost.

Windows 2003 Menu

Reject Data

Windows 2008 Menu

Reject

Description

If you reject the data, the node will be reverted to the snapshot that was taken when you selected the **Verify Group** or **Test Data** option. Any changes made on the node after that snapshot was created will be lost. This option essentially takes you back to where you were, allowing you the opportunity to check other nodes for more recent data.

GeoCluster Replicated Disk Status Resource

The function of the GeoCluster Replicated Disk Status resource (also displayed as GRD Status) varies between Windows 2003 and Windows 2008. In both operating systems, it is automatically created when the first GeoCluster Replicated Disk resource is created in a group. Once the status resource is created, it will exist as long as there is a GeoCluster Replicated Disk resource in the group. When the last GeoCluster Replicated Disk resource in a group is deleted, the status resource will be deleted. Only one status resource is created per group.

If you are using Windows 2003, the description of the status resource corresponds to various states of your Double-Take Availability data. By reviewing the status descriptions, you can tell at-a-glance the state of your Double-Take Availability data. If you are using Windows 2008, these status descriptions are seen directly in the GeoCluster Replicated Disk resource description, rather than the GeoCluster Replicated Disk Status resource.

For example, you may see the status "The status of all targets is OK." This indicates the data on each target node is in a good state. Another message may be "Target target_ name is queuing. Data in queue on target." This indicates the data on the specified target is not up-to-date. Because there is data in queue on the target, that has not been written to disk yet, the target data is out-of-date. Or you may see either of the following status descriptions.

- Target target_name is pending. Data integrity not guaranteed.
- Target target_name is suspect. Data integrity not guaranteed.

These messages indicate the data on the specified target node is not in a good state. This could be because a mirror is in progress, an operation has been dropped on the target, or another Double-Take Availabilityprocessing issue. Check the Double-Take Availability logs for more information. As long as the status is pending, data integrity cannot be guaranteed on the specified target node.

The text of the descriptions may vary between Windows 2003 and Windows 2008.

Another function of the status resource, for both Windows 2003 and Windows 2008, is to keep you from moving the GeoCluster Replicated Disk resource to another node at the wrong time and potentially corrupting your data. If the GeoCluster Replicated Disk resource was moved while the status resource is in a pending or queuing state, the new node would have incomplete data, which would then be replicated to the other nodes in the cluster. This resource safeguards you from corrupting your data.

Log files

Various Double-Take Availability components (Double-Take service, Replication Console, Failover Control Center, and the Command Line Client) generate a log file to gather alerts, which are notification, warning, and error messages. The log files are written to disk.

Each log file consists of a base name, a series number, and an extension.

- **Base Name**—The base name is determined by the application or process that is running.
  - **Double-Take Availability**—dtlog
  - **Replication Console**—mc
  - **Failover Control Center**—fcc
  - **Command Line Client**—dtcl
- **Series Number**—The series number ranges from 1 to 999. For example, Double-Take Availability begins logging messages to dtlog1. When this file reaches its maximum size, the next log file will be written to dtlog2. As long as log messages continue to be written, files dtlog3, dtlog4, dtlog5 will be opened and filled. When the maximum number of files is reached, which by default is 5, the oldest file is deleted when the sixth file is created. For example, when dtlog6 is created, dtlog1 is deleted and when dtlog7 is created, dtlog2 is deleted. When file dtlog999 is created and filled, dtlog1 will be re-created and Double-Take Availability will continue writing log messages to that file. In the event that a file cannot be removed, its number will be kept in the list, and on each successive file remove, the log writer will attempt to remove the oldest file in the list.
- **Extension**—The extension for each log file is .dtl.
  - **Double-Take Availability**—dtlog1.dtl, dtlog2.dtl
  - **Replication Console**—mc1.dtl, mc2.dtl
  - **Failover Control Center**—fcc1.dtl, fcc2.dtl
  - **Command Line Client**—dtcl1.dtl, dtcl2.dtl

Viewing the log file

You can view the Double-Take Availability log file through the Replication Console or through any text editor. You can view any of the other component log files through any text editor.

- Viewing the log file through the Replication Console
- Viewing the log file through a text editor

Viewing the log file through the Replication Console

1. Open a new message window using any of the following methods.
   - Right-click on the server that you want to monitor in the left pane and select **New**, **Message Window**.
   - Select the Message Window icon from the toolbar.
   - Select **Monitor**, **New Message Window** and identify the **Server** that you want to monitor.
2. Repeat step 1 if you want to open multiple message windows.

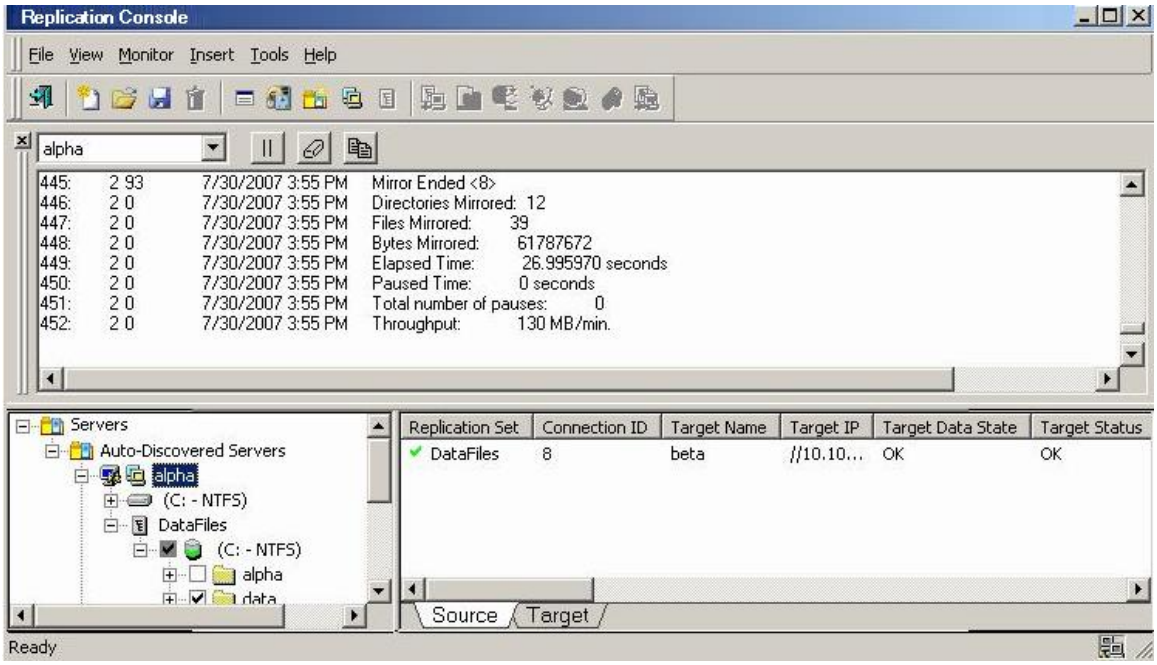

Note: The standard appearance of the message window is a white background. If your message window has a gray background, the window is inactive. The Replication Console may have lost communications with that server, for example, or you may no longer be logged into that server.

The message window is limited to the most recent 1000 lines. If any data is missing an entry in red will indicate the missing data. Regardless of the state of the message window, all data is maintained in the Double-Take Availability log on the server.

---

3. To control the window after it is created, use one of the following methods to access the control methods listed in the following table.
   - Right-click on the message window and select the appropriate control.
   - Select the appropriate toolbar control.
   - Select **Monitor**, the name of the message window, and the appropriate control.

Close 	Closes the message window
Clear 	Clears the message window
Pause/Resume 	Pauses and resumes the message window. Pausing prevents new messages from being displayed in the message window so that you are not returned to the bottom of the message window every time a new message arrives. The messages that occur while the window is logged are still logged to the Double-Take Availability log file. Resuming displays the messages that were held while the window was paused and continues to display any new messages. Pausing is automatically initiated if you scroll up in the message window. The display of new log messages will automatically resume when you scroll back to the bottom.
Copy 	Allows you to copy selected text
Options	This control is only available from the **Monitor** menu. Currently, there are no filter options available so this option only allows you to select a different server. In the future, this control will allow you to filter which messages to display.

4. To change which server you are viewing messages for, select a different machine from the drop down list on the toolbar. If necessary, the login process will be initiated.
5. To move the message window to other locations on your desktop, click and drag it to another area or double-click it to automatically undock it from the Replication Console.

Viewing the log file through a text editor

The log files can be viewed, from the location where Double-Take Availability is installed, with a standard text editor. The following list describes the information found in each column of the log file.

1. Date the message was generated
2. Time the message was generated
3. Process ID
4. Thread ID
5. Sequence number is an incremental counter that assigns a unique number to each message
6. The type or level of message displayed - 1 for warning or error message and 2 for informational message
7. Message ID
8. Message text

Sample Double-Take Availability log file

```
01/15/2010 14:14:18.3900 95 98 2 2 69 Kernel Started
01/15/2010 14:14:18.4200 95 98 3 2 10004 Valid Activation Key Detected :
01/15/2010 14:14:18.5350 98 170 4 2 52501 Target module loaded successfully
01/15/2010 14:14:18.6760 98 172 5 2 10004 Valid Activation Key Detected :
01/15/2010 14:14:18.9870 130 131 6 2 51501 Source module loaded successfully

01/15/2010 14:24:15.2070 130 132 7 2 72 Connection Request from
ip://206.31.4.305
01/15/2010 14:24:16.3090 131 133 8 2 600002 Unified login provides ADMIN
access
01/15/2010 14:24:40.9680 132 134 9 2 99 RepSet Modified: UserData
01/15/2010 14:25:22.4070 134 131 10 2 71 Originator Attempting
ip://206.31.4.305 01/15/2010 14:25:22.5030 134 131 11 2 0 Transmission
Create to ip://206.31.4.305.
01/15/2010 14:25:22.6060 135 133 12 2 500000 UserData is connected to
ip://206.31.4.305 01/15/2010 14:25:23.5030 136 98 13 2 87 Start Replication
on connection 1
```

Sample Replication Console log file

```
00/00/0000 00:00:00.0000 Application starting
09/11/2010 12:45:53.8980 704 1032 1 2 0 Could not find XML file: C:\Program
Files\Double-Take Software\Double-Take\Administrator.xml, default groups
will be added.
09/11/2010 12:45:53.9580 704 1032 2 2 0 Adding default group: Double-Take
```

```
Servers
09/11/2010 12:45:53.9580 704 1032 3 2 0 Adding default group: Double-Take
Servers\
Auto-Discovered Servers
09/11/2010 12:46:08.3390 704 1032 4 210004 Evaluation license expires in 95
day(s).
```

Filtering the log file

Log file output can be filtered using the LogViewer utility. Use the LogViewer command from the directory where Double-Take Availability is installed.

Command

LOGVIEWER

Description

The Double-Take Availability logging utility that filtersDouble-Take Availability log files

Syntax

LOGVIEWER [-PATH <*path*>] [-TYPE <*number*>] [-INCLUDE <*list*>] [-EXCLUDE <*list*>] [-NODATE] [-NOTIME] [-NOPID] [-NOTID] [-NOSEQ] [-NOTYPE] [-NOID] [-HELP]

Options

- PATH *path*—Specify the full path to the log file
- TYPE *number*—Allows you to filter the messages that are displayed. Specify 1 to display warning and error messages or specify 2 to display warnings, errors, and notifications
- INCLUDE—Only includes specified IDs. All other IDs will not be displayed in the output
- EXCLUDE—Excludes specified IDs. Ignore the specified IDs and display all others
- *list*—A comma-separated list of IDs or ID ranges that follows the INCLUDE and EXCLUDE switches. A space should separate the switch from the list but within the list, there should be no spaces. Ranges are specified with a begin and end number and separated with a dash (-).
- NODATE—Does not display the date in the output
- NOTIME—Does not display the time in the output
- NOPID—Does not display the process ID in the output
- NOTID—Does not display the thread ID in the output
- NOSEQ—Does not display the sequence number in the output
- NOTYPE—Does not display the message type number in the output

- NOID—Does not display the LogViewer ID in the output
- HELP—Displays the command options

Examples

- LogViewer -type 2
- LogViewer -include 200,400-500,10000-15000

Notes

The default setting is -type 2 which displays both type 1 and 2 messages.

Configuring the properties of the log file

1. To modify the maximum file size and the number of Double-Take Availability log files that are maintained, access the Server Properties dialog box by right-clicking a machine name in the left pane of the Replication Console and selecting **Properties**.
2. Select the **Logging** tab.
3. At the top of the window, **Folder** indicates the directory where the log files are located. The default is the directory where the Double-Take Availability program files are installed.
4. Modify any of the options under **Messages and Alerts**, if necessary.
   - **Maximum Length**—Specify the maximum length of the log file. The default size is 5242880 bytes and is limited by the available hard drive space.
   - **Maximum Files**—Specify the maximum number of log files that are maintained. The default is 5 and the maximum is 999.

   Note: If you change the **Maximum Length** or **Maximum Files**, you must restart the Double-Take service for the change to take effect.
5. Click **OK** to save the changes.

Double-Take Availability log messages

The following list describes some of the standard Double-Take Availability alerts that may be displayed in the log files.

Note: In this information, con_id refers to the unique connection ID assigned to each connection between a source replication set and a target.

There are several log messages with the ID of 0. See the description in the Message column in the log file.

---

7 Synchronous ioctl returned STATUS_PENDING

7 Failed to reset Replication Flags. Replication may not be performed correctly.

- Communication with the Double-Take Availability driver is not being performed correctly. A reboot is required to guarantee replication and data integrity.
- An error occurred between the Double-Take Availability driver and recent changes to the replication set. The possible resolutions are to undo the changes to the replication set, stop and restart Double-Take Availability, or reboot the server.

69 Double-Take kernel started on server_name

The Double-Take service was started on the Double-Take Availability server specified.

70 Double-Take kernel stopped

The Double-Take service was stopped on a Double-Take Availability server.

71 Originator attempting ip://xxx.xxx.xxx.xxx

A source is requesting to connect a replication set to a target machine.

72 Connection request from ip://xxx.xxx.xxx.xxx

A target machine has received a source machine's request to connect a replication set to the target.

73 Connected to ip://xxx.xxx.xxx.xxx

A source machine has successfully connected a replication set to a target machine.

74 Connection paused with ip://xxx.xxx.xxx.xxx

A network connection between the source and the target exists and is available for data transmission, but data is being held in queue and is not being transmitted to the target. This happens because the target machine cannot write data to disk fast enough. Double-Take Availability will resolve this issue on its own by transmitting the data in queue when the target catches up.

75 Connection resumed with ip://xxx.xxx.xxx.xxx

The transmission of data from the source machine to the target machine has resumed.

76 Connection failed to ip://xxx.xxx.xxx.xxx

An attempt to establish a network connection between a source machine and target machine has failed. Check your network connections and verify that the target machine is still online.

77 Connection lost with IP address address

The network connection previously established between a source machine and target machine has been lost. Check your network connections and troubleshoot to see why the connection was lost.

78 Auto-disconnect threshold has been reached.

The Double-Take Availability queue has exceeded its limit, and the auto-disconnect process will disconnect the source and target connection. The auto-reconnect process will automatically reestablish the connection if the auto-reconnect feature is enabled. If the auto-reconnect feature is not enabled, you must first verify that the connection between the source and target has been broken, and then manually reestablish the connection in the Replication Console.

79 Memory freed to bring Double-Take memory usage below the limit

Data in the source queue has been sent to the target machine, bringing the pagefile below its limit.

80 Trying to auto-retransmit to ip://xxx.xxx.xxx.xxx

Double-Take Availability is attempting to automatically reconnect previously established source and target connections after a server reboot or auto-disconnect. This is also referred to as the auto-reconnect process.

81 Schedule transmit start to target

A scheduled transmission of data from a source machine to a target machine has started. See the description in the Message column in the log file.

82 Schedule transmit end to target

A scheduled transmission of data from a source machine to a target machine has ended. See the description in the Message column in the log file.

85 repset has been auto-disconnected

Double-Take Availability automatically disconnects the source and target connection because the queue size has reached a specified size for this action.

87 Start replication on connection con_id

Data has started replicating from a source machine to a target machine.

88 Stop replication on connection con_id

Data has stopped replicating from a source machine to a target machine.

89 Mirror started con_id

Data is being mirrored from a source machine to a target machine.

90 Mirror stopped con_id

The process of mirroring data from a source machine to a target machine has stopped due to user intervention or an auto-disconnect. (This means the mirroring process was not completed.)

91 Mirror paused con_id

The process of mirroring data from a source machine to a target machine has paused because the target machine cannot write the data to disk fast enough. Double-Take Availability will resolve this issue on its own by transmitting the data in queue when the target catches up.

92 Mirror resumed con_id

The process of mirroring data from a source machine to a target machine has resumed.

93 Mirror ended con_id

The process of mirroring data from a source machine to a target machine has ended.

94 Verification started con_id

The verification process of confirming that the Double-Take Availability data on the target is identical to the data on the source has started.

95 Verification ended con_id

The verification process of confirming that the Double-Take Availability data on the target is identical to the data on the source has ended.

97 Restore started con_id

The restoration process of copying the up-to-date data from the target back to the original source machine has started.

98 Restore completed con_id

The restoration process of copying the up-to-date data from the target back to the original source machine has been completed.

99 RepSet Modified: repset_ name

This message means that the specified replication set has been modified.

100 Failover condition has been met and user intervention is required

Double-Take Availability has determined that the source has failed, and requires manual intervention to start the failover process.

101 Failover in progress!!!

The conditions for failover to occur have been met, and the failover process has started.

102 Target full!

The disk to which data is being written on the target is full. This issue may be resolved by deleting files on the target machine or by adding another disk.

801 Auto-disconnect has occurred on IP address with connection con_id Disconnected replication set name: repset_name.

Auto-disconnect has occurred for the specified connection. This is due to the source queue filling up because of a network or target failure or bottleneck.

10001 Activation key is not valid.

An invalid activation code was identified when the Double-Take service was started.

10002 Evaluation period has expired.

The evaluation license has expired.

10003 Activation code violation with machine machine_name

Duplicate single-server activation codes are being used on the servers, and Double-Take Availability is disabled.

10004 Valid activation key detected

A valid activation code was identified when the Double-Take service was started.

51001 Source module failed to load

The Double-Take Availability source module failed to load. Look at previous log messages to determine the reason. (Look for messages that indicate that either the activation code was invalid or the user-configurable source module was not set to load automatically at startup.) The source module may have been configured this way intentionally.

51501 Source module loaded successfully

The Double-Take Availability source module was loaded successfully.

51502 Source module already loaded

The Double-Take Availability source module was already loaded.

51503 Source module stopped

The Double-Take Availability source module stopped.

52000 The target has been paused due to manual intervention.

52000 The target has been resumed due to manual intervention

The target has been paused or resumed through user intervention.

52000 Unfinished Op error

This error message contains various Microsoft API codes. The text Code - <x> Internal <y> appears at the end of this message. The code value indicates why the operation failed, and the internal value indicates the type of operation that failed. These are the most common code values that appear in this error message.

- (5) Permission denied: The account running the Double-Take service does not have permission to update the file specified.
- (32) Sharing violation: Another application is using a particular file that Double-Take Availability is trying to update. Double-Take Availability will wait and try to update the file later.

- (112) Disk full: The disk to which data is being written on the target is full. This issue may be resolved by deleting files on the target machine or by adding another disk.

52501 Target module loaded successfully

The Double-Take Availability target module was loaded successfully.

52502 Target module already loaded

The Double-Take Availability target module was already loaded.

52503 Target module stopped

The Double-Take Availability target module stopped.

53001 File was missing from target

The verification process confirms that the files on the target are identical to the files on the source. This message would only appear if the verification process showed that a file on the source was missing from the target.

53003 Could not read filename

Double-Take Availability could not read a file on the source machine because the file may have been renamed or deleted. For example, temporary files show up in queue but do not show up during transmission. (No user action required.)

54000 Kernel started

The Double-Take service was started.

54001 Failover module failed to load

The Double-Take Availability failover module failed to load. Look at previous log messages to determine the reason.

54503 Failover module stopped

The Double-Take Availability failover module stopped.

99001 Starting source module low memory processing

The source's queue is full, and the auto-disconnect process will disconnect the source and target connection. The auto-reconnect process will automatically reestablish the connection if the auto-reconnect feature is enabled. If the auto-reconnect feature is not enabled, you must first verify that the connection between the source and target has been broken, and then manually reestablish the connection in the Replication Console.

99999 Application is terminating normally

The Double-Take service is shutting down normally.

503010 Asyncloctl for status thread 178 terminated, terminating the status thread

A Double-Take Availability process monitors the state of the Double-Take Availability driver. When the Double-Take service is shut down, the driver is shut down, and this process is terminated. (No user action required.)

600002 Unified login provides ADMIN access

600002 User has level access (x)

- Using the current login grants ADMIN access.
- The listed user has listed access level and access level ID.

700000 The source machine source_machine is not responding to a ping.

This occurs when all monitored IP addresses on the source machine stop responding to pings. Countdown to failover will begin at the first occurrence and will continue until the source machine responds or until failover occurs.

800000 Active Directory GetHostSpns function call failed

800000 Active Directory RemoveSpns function call failed

800000 Active Directory AddSpns function call failed

- Double-Take Availability failed to get the host SPN (Service Principal Name) from Active Directory.
- Double-Take Availability failed to remove an SPN from Active Directory.
- Double-Take Availability failed to add a host SPN to Active Directory.

Monitoring event messages

An event is a significant occurrence in the system or in an application that requires administrators to be notified. The operating system writes notifications for these events to a log that can be displayed using the Windows Event Viewer. Three different log files are generated: application, security, and system.

1. To access the Event Viewer, select **Programs**, **Administrative Tools**, **Event Viewer**.
2. Select the log to view (**System**, **Security**, or **Application**) from the left pane of the Event Viewer. The following information is displayed for an event in the right pane of the Event Viewer.
   - **Type**—A classification of the event, such as Error, Warning, Information, Success Audit, or Failure Audit.
   - **Date**—The date the event occurred.
   - **Time**—The time the event occurred.
   - **Source**—The software that logged the event, which can be either an application or a component of the system, such as a driver.
   - **Category**—A classification of the event.
   - **Event**—Shows an ID number to identify the specific event. The **Event** helps product-support representatives track events in the system.
   - **User**—Identifies the user that logged the event.
   - **Computer**—The name of the computer where the event occurred.
3. To view a detailed description, double-click an event. The additional information is displayed in the Event Properties screen.

Note: For additional information on customizing the Event Viewer (such as sorting the display, filtering the display, and so on), see your Windows reference guide or the Windows online help.

For a complete list of Double-Take events, see Event messages.

Event messages

The following table identifies the Double-Take Availability events.

---

1 This evaluation period has expired. Mirroring and replication have been stopped. To obtain a license, please contact your vendor.

Error—Contact your vendor to purchase either a single or site license.

2 The evaluation period expires in %1 day(s).

Information—Contact your vendor before the evaluation period expires to purchase either a single or site license.

3 The evaluation period has been activated and expires in %1 day(s).

Information—Contact your vendor before the evaluation period expires to purchase either a single or site license.

4 Duplicate activation codes detected on machine %1 from machine %2.

Warning—If you have an evaluation license or a site license, no action is necessary. If you have a single license, you must purchase either another single license or a site license.

5 This product edition can only be run on Windows Server or Advanced Server running the Server Appliance Kit.

Error—Verify your activation code has been entered correctly and contact technical support.

1000 An exception occurred: %1

Error—Run the installation and select Repair. Contact technical support if this event occurs again.

1001 The Double-Take counter DLL could not initialize the statistics handler object to gather performance data.

Error—Run the installation and select Repair. Contact technical support if this event occurs again.

1002 The Double-Take counter DLL could not map shared memory file containing the performance data.

Error—Run the installation and select Repair. Contact technical support if this event occurs again.

1003 The Double-Take counter DLL could not open the "Performance" key in the Double-Take section of the registry.

Error—Run the installation and select Repair. Contact technical support if this event occurs again.

1004 The Double-Take counter DLL could not read the "First Counter" value under the Double-Take\Performance Key.

Error—Run the installation and select Repair. Contact technical support if this event occurs again.

1005 The Double-Take counter DLL read the "First Help" value under the Double-Take\Performance Key.

Error—Run the installation and select Repair. Contact technical support if this event occurs again.

1006 The Double-Take counter DLL could not create event handler for the worker thread.

Error—Run the installation and select Repair. Contact technical support if this event occurs again.

3000 Logger service was successfully started.

Information—No action required.

3001 Logger service was successfully stopped.

Information—No action required.

4000 Kernel was successfully started.

Information—No action required.

4001 Target service was successfully started.

Information—No action required.

4002 Source service was successfully started.

Information—No action required.

4003 Source service was successfully stopped.

Information—No action required.

4004 Target service was successfully stopped.

Information—No action required.

4005 Kernel was successfully stopped.

Information—No action required.

4006 Service has aborted due to the following unrecoverable error: %1

Error—Restart the Double-Take service.

4007 Auto-disconnecting from %1 (%2) for Replication Set %3, ID: %4 due to %5

Warning—The connection is auto-disconnecting because the disk-based queue on the source has been filled, the service has encountered an unknown file ID, the target server has restarted, or an error has occurred during disk queuing on the source or target (for example, Double-Take Availability cannot read from or write to the transaction log file).

4008 Auto-disconnect has succeeded for %1 (%2) for Replication Set %3, ID: %4

Information—No action required.

4009 Auto-reconnecting Replication Set %1 to %2 (%3)

Information—No action required.

4010 Auto-reconnect has succeeded connecting Replication Set %1 to %2 (%3)

Information—No action required.

4011 Auto-reconnect has failed connecting Replication Set %1 to %2 (%3)

Error—Manually reestablish the replication set to target connection.

4012 %1

Warning—This is a placeholder message for many other messages. See the specific log message for additional details.

4013 %1

Information—This is a placeholder message for many other messages. See the specific log message for additional details.

4014 Service has started network transmission.

Information—No action required.

4015 Service has stopped network transmission.

Information—No action required.

4016 Service has established a connection to %1 (%2) for Replication Set %3, ID: %4

Information—No action required.

4017 Service has disconnected from %1 (%2) for Replication Set %3, ID: %4

Information—No action required.

4018 %1, however, mirroring and replication have been disabled as a restore is required due to a previous failover.

Warning—Perform a restoration.

4019 Service has started a mirror to %1 (%2) for Replication Set %3, ID: %4

Information—No action required.

4020 Service has paused a mirror to %1 (%2) for Replication Set %3, ID: %4

Information—No action required.

4021 Service has resumed a mirror to %1 (%2) for Replication Set %3, ID: %4

Information—No action required.

4022 Service has stopped a mirror to %1 for Replication Set %2, ID: %3, %4

Information—No action required.

4023 Service has completed a mirror to %1 %2 for Replication Set %3, ID: %4, %5

Information—No action required.

4024 Service has started Replication to %1 (%2) for Replication Set %3, ID: %4

Information—No action required.

4025 Service has stopped Replication to %1 (%2) for Replication Set %3, ID: %4

Information—No action required.

4026 The target has been paused due to user intervention.

Information—No action required.

4027 The target has been resumed due to user intervention.

Information—No action required.

4028 Registration of service class with Active Directory failed. Verify that the Active Directory server is up and the service has the proper permissions to update its entries.

Warning—Verify that the Active Directory server is running and that the Double-Take service has permission to update Active Directory.

4029 Registration of service instance with Active Directory failed. Verify that the Active Directory server is up and the service has the proper permissions to update its entries.

Warning—Verify that the Active Directory server is running and that the Double-Take service has permission to update Active Directory.

4030 RSResource.dll has an unknown error. The product functionality has been disabled.

Error—Reinstall the software, using the installation Repair option, to install a new copy of the RSResource.dll. Contact technical support if this error persists.

4031 RSResource.dll could not be opened. The product functionality has been disabled.

Error—Reinstall the software, using the installation Repair option, to install a new copy of the RSResource.dll. Contact technical support if this error persists.

4032 The RSResource.dll component version does not match the component version expected by the product. The product functionality has been disabled.

Error—Reinstall the software, using the installation Repair option, to install a new copy of the RSResource.dll. Contact technical support if this error persists.

4033 RSResource.dll build version is invalid. The product functionality has been disabled.

Error—Reinstall the software, using the installation Repair option, to install a new copy of the RSResource.dll. Contact technical support if this error persists.

4034 Error verifying the service name. The product functionality has been disabled.

Error—Reinstall the software, using the installation Repair option, to install a new copy of the RSResource.dll. Contact technical support if this error persists.

4035 Error verifying the product name. The product functionality has been disabled.

Error—Reinstall the software, using the installation Repair option, to install a new copy of the RSResource.dll. Contact technical support if this error persists.

4036 Error verifying the vendor name. The product functionality has been disabled.

Error—Reinstall the software, using the installation Repair option, to install a new copy of the RSResource.dll. Contact technical support if this error persists.

4037 Error verifying the vendor URL name. The product functionality has been disabled.

Error—Reinstall the software, using the installation Repair option, to install a new copy of the RSResource.dll. Contact technical support if this error persists.

4038 Error verifying the product code. The product functionality has been disabled.

Error—Reinstall the software, using the installation Repair option, to install a new copy of the RSResource.dll. Contact technical support if this error persists.

4039 Error while reading RSResource.dll. The product functionality has been disabled.

Error—Reinstall the software, using the installation Repair option, to install a new copy of the RSResource.dll. Contact technical support if this error persists.

4040 The product code is illegal for this computer hardware. The product functionality has been disabled.

Error—Reinstall the software, using the installation Repair option, to install a new copy of the RSResource.dll. Contact technical support if this error persists.

4041 The product code is illegal for this operating system version. The product functionality has been disabled.

Error—Reinstall the software, using the installation Repair option, to install a new copy of the RSResource.dll. Contact technical support if this error persists.

4042 The product code requires installing the Windows Server Appliance Kit. The product functionality has been disabled.

Error—Reinstall the software, using the installation Repair option, to install a new copy of the RSResource.dll. Contact technical support if this error persists.

4043 This product can only be run on a limited number of processors and this server exceeds the limit. The product functionality has been disabled.

Error—Reinstall the software, using the installation Repair option, to install a new copy of the RSResource.dll. Contact technical support if this error persists.

4044 An error was encountered and replication has been stopped. It is necessary to stop and restart the service to correct this error.

Error—Contact technical support if this error persists.

4045 %1 value must be between 1025 and 65535. Using default of %2.

Error—Verify that the Double-Take Availability port value you are trying to use is within the valid range. If it is not, it will automatically be reset to the default value.

4046 This service failed to start because of a possible port conflict. Win32 error: %1

Error—Verify that the Double-Take Availability ports are not conflicting with ports used by other applications.

4047 Could not load ZLIB DLL %1. Some levels of compression will not be available.

Error—The compression levels available depend on your operating system. You can reinstall the software, using the installation Repair option, to install a new copy of the DynaZip.dll, or contact technical support if this error persists.

4048 Service has started a delete orphans task to %1 (%2) for Replication Set %3, ID: %4

Information—No action required.

4049 Service has paused a delete orphans task to %1 (%2) for Replication Set %3, ID: %4

Information—No action required.

4050 Service has resumed a delete orphans task to %1 (%2) for Replication Set %3, ID: %4

Information—No action required.

4051 Service has stopped a delete orphans task to %1 (%2) for Replication Set %3, ID: %4

Information—No action required.

4052 Service has completed a delete orphans task to %1 (%2) for Replication Set %3, ID: %4

Information—No action required.

4053 Service has started a restore task to %1 (%2) for Replication Set %3, ID: %4

Information—No action required.

4054 Service has paused a restore task to %1 (%2) for Replication Set %3, ID: %4

Information—No action required.

4055 Service has resumed a restore task to %1 (%2) for Replication Set %3, ID: %4

Information—No action required.

4056 Service has stopped a restore task to %1 (%2) for Replication Set %3, ID: %4

Information—No action required.

4057 Service has completed a restore task to %1 (%2) for Replication Set %3, ID: %4

Information—No action required.

4058 Service has started a verification task to %1 (%2) for Replication Set %3, ID: %4

Information—No action required.

4059 Service has paused a verification task to %1 (%2) for Replication Set %3, ID: %4

Information—No action required.

4060 Service has resumed a verification task to %1 (%2) for Replication Set %3, ID: %4

Information—No action required.

4061 Service has stopped a verification task to %1 (%2) for Replication Set %3, ID: %4

Information—No action required.

4062 Service has completed a verification task to %1 (%2) for Replication Set %3, ID: %4

Information—No action required.

4063 Bandwidth limit to %1 (%2) has changed to %3.

Information—No action required.

4064 Bandwidth limit to %1 (%2) is now in the "%3" period at %4.

Information—No action required.

4065 Target data state for connection %1 from %2 (%3) has changed because %4.

Warning—No action required.

4066 The product code requires a virtual server environment. The product functionality has been disabled.

Error—The activation code you are using is for the Virtual SystemsTM edition. This code will not work on non-virtual server environments.

4067 No replication ops have been received from the driver for an extended period of time.

Error—Check other messages for errors with the Double-Take Availability drivers, and correct as required. If there are no driver messages, verify that your drives are connected to the source. If this error persists, contact technical support.

4068 Failed to write to a replicating volume.

Error—Reboot the source server. Contact technical support if this event occurs again.

4069 The option MoveOrphansDir has been updated because it was missing or empty.

Warning—No action required.

4070 An error occurred while reading data for connection %1. All data needs to be remirrored. See the log for details.

Error—Double-Take Availability will automatically remirror the data. See the Double-Take Availability log file for details.

4096 The registry parameter %2 is unknown.

Warning—Delete the parameter and report this issue to technical support.

4097 Failed to initialize WMI support. The last Word in the Data Window is the NT status code.

Warning—No action required.

4097 The file system filter failed to load. Replication will not occur. Reboot your server and contact technical support if this error occurs again. The last Word in the Data window is the NT status code.

Error—Reboot your server and contact technical support if this event occurs again.

4098 The registry parameters failed to load, so the default configuration values will be used. The last Word in the Data window is the NT status code.

Warning—No action required.

4098 The control device %2 was not created. Communication with the service will be disabled. Reboot the server and contact technical support if this error occurs again. The last Word in the Data window is the NT status code.

Error—Reboot your server and contact technical support if this event occurs again.

4099 The driver detected a hard link for a file on drive %2. Hard links are not supported. Changes to this file will not be replicated.

Warning—Hard links are not supported.

4099 The driver failed to register with filter manager. Reboot the server and contact technical support if this error occurs again. The last Word in the Data window is the NT status code.

Error—Reboot your server and contact technical support if this event occurs again.

4100 Product activation code is invalid. Please check that it is typed correctly and is valid for the version of the operating system in use.

Error—If you are in the process of installing Double-Take Availability, verify that you are using a 24 character alpha-numeric code. If Double-Take Availability is already installed, confirm that the code entered is correct. If the code appears to be correct, contact technical support.

4100 The versions of the driver and the filter driver do not match. Replication will not occur. Reboot your server. If this error occurs again, reinstall the software. Contact technical support if this error occurs after the software has been reinstalled. The last three Words in the Data window are the NT status code and the driver version numbers.

Error—Reboot your server. Reinstall the software if this event occurs again. Contact technical support if this event occurs after reinstalling the software.

4101 This service will not run on this device. Contact your sales representative for upgrade procedures.

Error—The activation code does not match the type of server you are attempting to run on. Contact your vendor for a new activation code or contact technical support.

4110 Target cannot write %1 due to target disk being full. Operation will be retried (%2 times or forever)

Warning—The disk on the target is full. The operation will be retried according to the TGExecutionRetryLimit setting.

4111 Target can not write %1 due to a sharing violation. Operation will be retried (%2 times or forever)

Warning—A sharing violation error is prohibiting Double-Take Availability from writing on the target. The operation will be retried according to the TGExecutionRetryLimit setting.

4112 Target can not write %1 due to access denied. Operation will be retried (%2 times or forever)

Warning—An access denied error is prohibiting Double-Take Availability from writing on the target. The operation will be retried according to the TGExecutionRetryLimit setting.

4113 Target can not write %1 due to an unknown reason. Operation will be retried (%2 times or forever). Please check the log files for further information on the error.

Warning—An unknown error is prohibiting Double-Take Availability from writing on the target. The operation will be retried according to the TGExecutionRetryLimit setting.

4120 Target write to %1 was completed successfully after %2 retries.

Information—No action required.

4150 Target write %1 failed after %2 retries and will be discarded. See the event log or log files for error conditions. After correcting the problem, you should re-mirror or run a verify to resynchronize the changes.

Error—The operation has been retried according to the TGExecutionRetryLimit setting but was not able to be written to the target and the operation was discarded. Correct the problem and remirror the files.

4155 The service was unable to complete a file system operation in the allotted time. See the log files for error conditions. After correcting the problem, remirror or perform a verification with remirror to synchronize the changes.

Warning—Correct the file system error and then remirror or perform a verification with remirror to synchronize the changes.

4200 In band task %1 submitted from %2 by %3 at %4

Information—No action required.

4201 In band task %1 discarded (submitted from %2 by %3 at %4)

Warning—A task may be discarded in the following scenarios: all connections to a target are manually disconnected, replication is stopped for all connections to a target, or an auto-disconnect occurs. If one of these

scenarios did not cause the task to be discarded, contact technical support.

4202 Running %1 in band script: %2 (task %3 submitted from %4 by %5 at %6)

Information—No action required.

4203 Completed run of in band script: %1 (exit code %2)

Information—No action required.

4204 Error running in band script: %1

Error—Review the task and its associated script(s) for syntax errors.

4205 Timeout (%1 seconds) running in band script: %2

Warning—The timeout specified for the script to complete has expired. Normal processing will continue. You may need to manually terminate the script if it will never complete.

4206 Run timeout disabled for in band script: %1

Warning—The timeout period was set to zero (0). Double-Take Availability will not wait for the script to complete before continuing. No action is required.

4207 In band scripts disabled by server - no attempt will be made to run %1

Warning—Enable task command processing.

4300 A connection request was received on the target before the persistent target paths could be loaded.

Error—You may need to disconnect and reconnect your replication set.

4301 Unable to block target paths, the driver is unavailable.

Error—If you need to block your target paths, contact technical support.

4302 Target Path %1 has been successfully blocked

Information—No action required.

4303 Blocking of target path: %1 failed. Error Code: %2

Warning—If you need to block your target paths, contact technical support.

4304 Target Path %1 has been successfully unblocked

Information—No action required.

4305 Unblocking of target path: %1 failed. Error Code: %2

Warning—If you need to unblock your target paths, contact technical support.

4306 Target paths for source %1 (%2) Connection id: %3 are already blocked

Warning—No action required.

4307 Target paths for source %1 (%2) Connection id: %3 are already unblocked

Warning—No action required.

4308 Error loading target paths for blocking, registry key %1 has been corrupted.

Error—If you need to block your target paths, contact technical support.

4400 Failed to create snapshot set for source %1 (%2) Connection ID: %3. Error: %4

Error—The snapshot could not be created. This may be due to a lack of disk space or memory or another reason. The error code is the Microsoft VSS error. Check your VSS documentation or contact technical support.

4401 Failed to delete automatic snapshot set for source %1 (%2) Connection ID: %3. Error: %4

Error—The automatic snapshot could not be deleted. This may be due to a lack of memory, the file does not exist, or another reason. The error code is the Microsoft Volume Shadow Copy error. Check your Volume Shadow Copy documentation or contact technical support.

4402 Failed to delete snapshot set for source %1 (%2) Connection ID: %3. Error: %4

Error—The snapshot could not be deleted. This may be due to a lack of memory, the file does not exist, or another reason. The error code is the Microsoft Volume Shadow Copy error. Check your Volume Shadow Copy documentation or contact technical support.

4403 A scheduled snapshot could not be created for source %1 (%2) Connection ID: %3. because the target data was in a bad state. A snapshot will automatically be created when the target data reaches a good state.

Error—No action required. A snapshot will automatically be created when the target data reaches a good state.

4404 Set snapshot schedule for source %1 (%2) connection %3 to every %4 minutes. Next snapshot: %5.

Information—No action required.

4405 Removed snapshot schedule for source %1 (%2) connection %3.

Information—No action required.

4406 Enabled snapshot schedule for source %1 (%2) connection %3.

Information—No action required.

4407 Disabled snapshot schedule for source %1 (%2) connection %3.

Information—No action required.

4408 %1 was unable to move some orphans for source %2 on connection ID %3. Check the %1 logs for further details.

Warning—Orphan files could not be moved. For example, the location could be out of disk space. Check the Double-Take Availability log for more information.

4409 %3 was unable to delete some orphans for source %1 on connection ID %2. Check the %3 logs for further details.

Warning—Orphan files could not be deleted. Check the Double-Take Availability log for more information.

4410 The registry hive dump failed with an of error: %1.

Error—Contact technical support.

4411 The Service has detected that port %1 is being %2 by the Windows Firewall.

Warning—The firewall port needs to be unblocked or restrictions against Double-Take Availability removed so that Double-Take Availability data can be transmitted.

5000 Server Monitor service was successfully started.

Information—No action required.

5001 Server Monitor service was successfully stopped.

Information—No action required.

5002 Placeholders were modified to %1.

Information—No action required.

5100 Failover completed for %1.

Information—No action required.

5101 IP address %1 with subnet mask %2 was added to target machine's %3 adapter.

Information—No action required.

5102 %1 has reached a failover condition. A response from the user is required before failover can take place.

Warning—User intervention has been configured. Open the Failover Control Center and accept or decline the failover prompt.

5103 Started adding drive shares from %1 to %2.

Information—No action required.

5104 %1 drive shares were taken over by %2.

Information—No action required.

5105 Attempting to run the %1 script.

Information—No action required.

5106 The %1 script ran successfully.

Information—No action required.

5107 Error occurred in running %1 script.

Error—Verify that the script identified exists with the proper permissions.

5108 The source machine %1 is not responding to a ping.

Error—This occurs when all monitored IP addresses on the source machine stop responding to pings. Countdown to failover will begin at the first occurrence and will continue until the source machine responds or until failover occurs.

5109 The public NIC on source machine %1 is not responding to a ping.

Error—The failover target did not receive an answer to its ping of the source machine. Eventually, a failover will result. Investigate possible errors (down server, network error, etc.).

5110 The %1 script "%2" is still running.

Information—No action required.

5200 Failback completed for %1.

Information—No action required.

5201 IP address %1 was removed from target machine's %2 adapter.

Information—No action required.

5202 Unable to Failback properly because IP address %1 was missing a corresponding SubNet Mask.

Error—Contact technical support.

5300 The following IP address was added to target's monitoring list: %1

Information—No action required.

5301 The following IP address was removed from target's monitoring list: %1

Information—No action required.

5302 Drive share information for %1 has been updated on the target machine.

Information—No action required.

5303 The application monitor script has started successfully.

Information—No action required.

5304 The application monitor script has finished successfully.

Information—No action required.

5305 The application monitor has found the %1 service stopped.

Warning—Double-Take Availability Application Manager will attempt to restart the service.

5306 The application monitor has restarted the %1 service.

Warning—No action required.

5307 The application monitor cannot contact the server %1.

Error—Verify the server is running. Verify available network communications with the server.

5400 Broadcasted new MAC address %1 for IP address %2.

Information—No action required.

5500 Could not connect to e-mail server. Check to make sure the SMTP server %1 is available (error code: %2).

Warning—Double-Take Availability could not connect to your SMTP server or the username and/or password supplied is incorrect. Verify that SMTP server is available and that you have identified it correctly in your e-mail notification configuration. Also verify that your username and password have been entered correctly.

5501 E-mail notification could not be enabled (error code: %1).

Warning—This alert occurs if there is an unexpected error enabling e-mail notification during service startup. Check to see if any other errors related to e-mail notification have been logged. Also, check to make sure the Windows Management Instrumentation (WMI) service is enabled. If neither of these apply, contact technical support.

5502 E-mail notification could not be initialized. Check to make sure Internet Explorer 5.0 or later is installed.

Warning—E-mail notification no longer requires Internet Explorer 5.0 or later. If you receive this error, contact technical support.

5503 E-mail notification could not be processed. Check to make sure the correct version of SMTPMail.DLL is registered on the system (error code: %1).

Warning—If you are using Double-Take Availability 4.4.2.1 or earlier and Windows NT 4.0, e-mail notification requires Windows Management Instrumentation (WMI) to be installed. Verify that you have it installed on the Double-Take Availability server.

5504 Could not load LocalRS.dll (for e-mail notification).

Warning—This alert occurs if there is an error loading the resource DLL for the service. Typically, this is caused by a missing LocalRS.dll file. Reinstall the software, using the installation Repair option, to install a new copy of the LocalRS.dll. Contact technical support if this error persists.

5505 E-mail could not be sent. Check e-mail settings (error code: %1).

Warning—Verify that the e-mail server that you have identified in your e-mail notification configuration is correct.

5506 One or more required e-mail settings have not been specified (error code: %1).

Warning—At a minimum, you must specify the e-mail server, the From and To addresses, and at least one type of event to include.

5507 E-mail notification could not be initialized. Check to make sure WMI is installed and available (error code: %1).

Warning—If you are using Double-Take Availability 4.4.2.1 or earlier and Windows NT 4.0, e-mail notification requires Windows Management Instrumentation (WMI) to be installed. Verify that you have it installed on the Double-Take Availability server.

5508 An error occurred connecting to the WMI namespace. Check to make sure the Windows Management Instrumentation service is not disabled (error code %1).

Warning—This alert occurs if there is an error with the Windows Management Instrumentation (WMI) service. Verify that you have it installed on the Double-Take Availability server and that it is enabled.

5600 Part or all of the e-mail setting %1 is not in a valid format.

Warning—Verify that the include categories and exclude ID list are identified and formatted correctly.

7106 The driver was unable to get valid name information from the Filter Manager for the file %2. (Filename may be truncated.) It cannot be replicated. Please contact technical support.

Error—Contact technical support.

7107 The driver was unable to get valid name information from the Filter Manager for a file. It cannot be replicated. Please contact technical support.

Error—Contact technical support.

8100 The driver encountered an unrecoverable internal error. Contact technical support. The last Word in the Data window is the internal error code.

Error—Contact technical support.

8192 Driver failed to allocate Kernel memory. Replication is stopped and server must be rebooted for replication to continue. The last word in the data window is the tag of the allocation that failed.

Error—Reboot the server and contact technical support if this event occurs again.

8192 Kernel memory is exhausted. Replication is stopped. This may have been caused by low system resources.

Error—Reboot the server and contact technical support if this event occurs again.

8193 The driver failed to create a thread required for normal operation. This may have been caused by low system resources. Reboot your server and contact technical support if this error occurs again. The last Word in the Data window is the NT status code.

Error—Reboot the server and contact technical support if this event occurs again.

8196 The maximum amount of memory for replication queuing has been reached. Replication is stopped and memory is being freed.

Warning—Contact technical support if this event occurs again.

8198 The driver registry path could not be saved. The default registry path will be used.

Warning—No action required.

8200 The driver failed to allocate a buffer for a file name longer than 260 characters. The file will be skipped. The last Word in the Data window is the NT status code.

Warning—Reboot the server and contact technical support if this event occurs again.

9000 The driver has failed to process a rename operation. The driver will resend the rename operation. This message is only a warning. If you receive this message

repeatedly, contact technical support. The last Word in the Data window is the NT status code.

Warning—Contact technical support if this event occurs again.

9100 The driver encountered an error opening a file from the service. Check the Event Viewer Application log for additional service information or contact technical support. The last Word in the Data window is the exception code.

Error—Check for related service messages. Contact technical support if this event occurs again.

9101 The driver encountered an error reading from the service input buffer. Check the Event Viewer Application log for additional service information or contact technical support. The last Word in the Data window is the exception code.

Error—Check for related service messages. Contact technical support if this event occurs again.

9102 The driver encountered an error writing to the service output buffer. Check the Event Viewer Application log for additional service information or contact technical support. The last Word in the Data window is the exception code.

Error—Check for related service messages. Contact technical support if this event occurs again.

9103 The driver encountered an error writing to the service input buffer. Check the Event Viewer Application log for additional service information or contact technical support. The last Word in the Data window is the exception code.

Error—Check for related service messages. Contact technical support if this event occurs again.

9104 The driver encountered an error querying for file security from the service input buffer. Check the Event Viewer Application log for additional service information or contact technical support. The last Word in the Data window is the exception code.

Error—Check for related service messages. Contact technical support if this event occurs again.

9105 The driver encountered an error querying for file security from the service input buffer. Check the Event Viewer Application log for additional service information or contact technical support. The last Word in the Data window is the exception code.

Error—Check for related service messages. Contact technical support if this event occurs again.

9106 The driver encountered an error writing file security data to the service input buffer. Check the Event Viewer Application log for additional service information or contact technical support. The last Word in the Data window is the exception code.

Error—Check for related service messages. Contact technical support if this event occurs again.

9107 The driver encountered an error querying for an allocated range from the service input buffer. Check the Event Viewer Application log for additional service information or contact technical support. The last Word in the Data window is the exception code.

Error—Check for related service messages. Contact technical support if this event occurs again.

9108 The driver encountered an error querying for an allocated range from the service output buffer. Check the Event Viewer Application log for additional service information or contact technical support. The last Word in the Data window is the exception code.

Error—Check for related service messages. Contact technical support if this event occurs again.

9109 The driver encountered an error writing an allocated range to the service input buffer. Check the Event Viewer Application log for additional service information or contact technical support. The last Word in the Data window is the exception code.

Error—Check for related service messages. Contact technical support if this event occurs again.

9110 The driver encountered an error querying for a directory from the service input buffer. Check the Event Viewer Application log for additional service information or contact technical support. The last Word in the Data window is the exception code.

Error—Check for related service messages. Contact technical support if this event occurs again.

9111 The driver encountered an error querying for a directory from the service output buffer. Check the Event Viewer Application log for additional service information or contact technical support. The last Word in the Data window is the exception code.

Error—Check for related service messages. Contact technical support if this event occurs again.

9112 The driver encountered an error writing a directory query to the service input buffer. Check the Event Viewer Application log for additional service information or contact technical support. The last Word in the Data window is the exception code.

Error—Check for related service messages. Contact technical support if this event occurs again.

9113 The driver encountered an error querying a stream from the service input buffer. Check the Event Viewer Application log for additional service information or contact technical support. The last Word in the Data window is the exception code.

Error—Check for related service messages. Contact technical support if this event occurs again.

9114 The driver encountered an error writing a stream query to the service output buffer. Check the Event Viewer Application log for additional service information or contact technical support. The last Word in the Data window is the exception code.

Error—Check for related service messages. Contact technical support if this event occurs again.

9115 The driver encountered an error writing a stream query to the service output buffer. Check the Event Viewer Application log for additional service information or contact technical support. The last Word in the Data window is the exception code.

Error—Check for related service messages. Contact technical support if this event occurs again.

9116 The driver has failed to close a file handle. If you receive this message repeatedly, contact technical support. The last Word in the Data window is the NT status code.

Error—Contact technical support.

9117 The driver encountered an error querying for extended attributes from the service input buffer. Check the Event Viewer Application log for additional service information or contact technical support. The last Word in the Data window is the exception code.

Error—Check for related service messages. Contact technical support if this event occurs again.

9118 The driver encountered an error writing extended attributes to the service output buffer. Check the Event Viewer Application log for additional service

information or contact technical support. The last Word in the Data window is the exception code.

Error—Check for related service messages. Contact technical support if this event occurs again.

9119 The driver encountered an error writing extended attributes status to the service input buffer. Check the Event Viewer Application log for additional service information or contact technical support. The last Word in the Data window is the exception code.

Error—Check for related service messages. Contact technical support if this event occurs again.

9120 The driver encountered an error querying for file information from the service input buffer. Check the Event Viewer Application log for additional service information or contact technical support. The last Word in the Data window is the exception code.

Error—Check for related service messages. Contact technical support if this event occurs again.

9121 The driver encountered an error writing file information to the service output buffer. Check the Event Viewer Application log for additional service information or contact technical support. The last Word in the Data window is the exception code.

Error—Check for related service messages. Contact technical support if this event occurs again.

9122 The driver encountered an error writing file information status to the service input buffer. Check the Event Viewer Application log for additional service information or contact technical support. The last Word in the Data window is the exception code.

Error—Check for related service messages. Contact technical support if this event occurs again.

9123 The driver encountered an error querying for fsctl information from the service input buffer. Check the Event Viewer Application log for additional service information or contact technical support. The last Word in the Data window is the exception code.

Error—Check for related service messages. Contact technical support if this event occurs again.

9124 The driver encountered an error writing fsctl information to the service output buffer. Check the Event Viewer Application log for additional service information

or contact technical support. The last Word in the Data window is the exception code.

Error—Check for related service messages. Contact technical support if this event occurs again.

9125 The driver encountered an error writing fsctl status to the service input buffer. Check the Event Viewer Application log for additional service information or contact technical support. The last Word in the Data window is the exception code.

Error—Check for related service messages. Contact technical support if this event occurs again.

10000 This message is only a placeholder warning. The last Word in the Data window is the NT status code.

Warning—No action required.

10000 Connect failed to node %1 for resource %2. Adding node to reconnect list.

Error—Ensure that GeoCluster is running on all possible owners and that it can communicate on the network selected for mirroring and replication traffic. GeoCluster will try to reestablish a connection using the check unresponsive node interval specified for the resource.

10001 Reconnect succeeded to node %1 for resource %2. Will be added as a possible owner when mirror is complete.

Information—No action required.

10002 Disk check failed on node %1 for resource %2. Removing as a possible owner.

Error—Ensure that GeoCluster is running on all possible owners and that it can communicate on the public network. Also ensure that the disk specified for the resource is functioning correctly on all possible owners.

10003 Owner %1 of the quorum resource %2 couldn't access the arbitration path %3. Network may be down.

Error—Ensure that the network used to access the arbitration path is up and that the server is operational. Also ensure that the arbitration share path does exist and that the account running the cluster service has write privileges to the share path.

10004 Failover of the group %1 is being delayed. Group will be brought online when the target queue is below the limit or the timeout has expired.

Warning—No action required.

10005 Node %1 is taking ownership of the group %2. The group will be brought online on this node.

Information—No action required.

10006 The cluster notification thread failed to start on node %1 for resource %2. The resource should be taken offline and brought back online.

Warning—Take the resource offline and bring it back online.

10007 The user %1 has reverted a snapshot for the %2 resource on node %3.

Warning—No action required. The snapshot you selected will be reverted.

10008 The user %1 has discarded queued data for the %2 resource on node %3.

Warning—No action required. The queue you selected will be discarded.

10009 The user %1 is verifying data for the %2 resource on node %3.

Warning—A snapshot of the current data has been taken. After you have verified the data, accept or reject the data.

10010 The user %1 has rejected the data for the %2 resource on node %3.

Warning—No action required. Since the data was rejected, the data has been reverted to the snapshot taken when the data was selected for verification.

10011 The user %1 has accepted the data for the %2 resource on node %3.

Warning—No action required. The current data will be used.

10012 The GeoCluster Replicated Disk resource %1 has been set to validate its data. No data replication is occurring to the remaining nodes in the cluster. Please Accept or Reject the data by right-clicking on the resource and selecting the appropriate option.

Warning—Replication has been stopped because of the validation request. Accept or reject the data on the node by right-clicking on the resource and selecting the appropriate option.

10100 The driver could not recall a file because it did not have a token for impersonation. The security provider service should set this token. The last Word in the Data window is the exception code.

Error—Contact technical support if this event occurs again.

10101 The driver could not access the file in the archive bin, due to a failed impersonation attempt. The last Word in the Data window is the exception code.

Error—Contact technical support if this event occurs again.

10102 The driver could not recall the file. The last Word in the Data window is the exception code.

Error—Contact technical support if this event occurs again.

11000 Service has started an archive to %1 (%2) for Replication Set %3, ID: %4

Information—No action required.

11001 Service has completed an archive to %1 (%2) for Replication Set %3, ID: %4, %5

Information—No action required.

11002 Service has started a recall from %1 (%2) for Replication Set %3, ID: %4

Information—No action required.

11003 Service has completed a recall from %1 (%2) for Replication Set %3, ID: %4, %5

Information—No action required.

11004 Service has failed connection to the RepHSM driver. %1

Warning—Reboot the server or manually restart the RepHSM.sys driver.

11005 Service has aborted the archive operation.

Warning—Verify the activation code on the source and target is valid for archiving. Reboot an unlicensed server.

11006 Service has aborted the archive recall operation.

Warning—Verify the activation code on the source and target is valid for archiving. Reboot an unlicensed server.

11007 Verification has finished with errors. %1

Warning—Review the verification log to correct or accept the errors.

11008 Archive feature is not supported on volume %1

Warning—The source and target must be NTFS for archiving functionality

11009 Service has started an archive preview to %1 (%2) for Replication Set %3, ID: %4

Information—No action required.

11010 Service has completed an archive preview to %1 (%2) for Replication Set %3, ID: %4

Information—No action required.

11011 Service has aborted the archive preview operation.

Warning—Verify the activation code on the source and target is valid for archiving. Reboot an unlicensed server.

12000 The service has started.

Information—This message refers to the Double-Take Recall service. No action required.

12001 The service failed to start.

Error—Check the user name and password for the Double-Take Recall service to ensure validity. Reinstall the software if this event occurs again.

12002 The service has stopped.

Information—This message indicates a system shutdown or the user stopped the Double-Take Recall service. No action is required.

12003 The service failed to create a stop control event. (Error %1)

Error—Restart the Double-Take Recall service. Reinstall the software if this event occurs again.

12004 RegisterServiceCtrlHandler failed. (Error %1)

Error—Restart the Double-Take Recall service. Reinstall the software if this event occurs again.

12005 Service encountered SetServiceStatus error (Error %1)

Error—Restart the Double-Take Recall service. Reinstall the software if this event occurs again.

12006 Service could not get handle to driver for security update. (Error %1)

Error—The Double-Take Recall service could not connect to the Double-Take Recall archiving driver. Reboot the server and reinstall the software if this event occurs again.

12007 Service failed a periodic security update. (Error %1)

Warning—This message refers to the Double-Take Recall service. The operation will be performed every five minutes. Reinstall the software if this event occurs after five minutes.

12288 The driver encountered an error accessing a buffer from the service. Contact technical support. The last Word in the Data window is the exception code.

Error—Contact technical support.

16384 The driver encountered an unrecoverable error. Contact technical support.

Error—Contact technical support

16385 The driver encountered an unexpected internal result. Contact technical support. The last Word in the Data window is the NT status code.

Error—Contact technical support.

16393 The driver encountered an internal error. Contact technical support. The last Word in the Data window is the internal error code.

Error—Contact technical support.

16395 The driver detected a memory error which may have been caused by a bad driver or faulty hardware. Contact technical support. The last Word in the Data window is the internal error code.

Error—Contact technical support.

16396 The driver failed to create work queues for normal operation. This may have been caused by low system resources. Reboot the server and contact technical support if this error occurs again. The last Word in the Data window is the NT status code.

Error—Reboot the server and contact technical support if this event occurs again.

16400 RepDrv has encountered an unexpected condition, usually caused by low kernel memory. Unless otherwise mentioned, this event has already been handled and your data remains protected. If you continue to receive these events or have further questions please contact tech support.

Information—No action required.

E-mailing event messages

You can e-mail Double-Take Availability event messages to specific addresses. The subject of the e-mail will contain an optional prefix, the server name where the message was logged, the message ID, and the severity level (information, warning, or error). The text of the message will be displayed in the body of the e-mail message.

1. Open the Replication Console and right-click the server on the left pane of the Replication Console.
2. Select **Properties**.
3. Select the **E-mail Notification** tab.

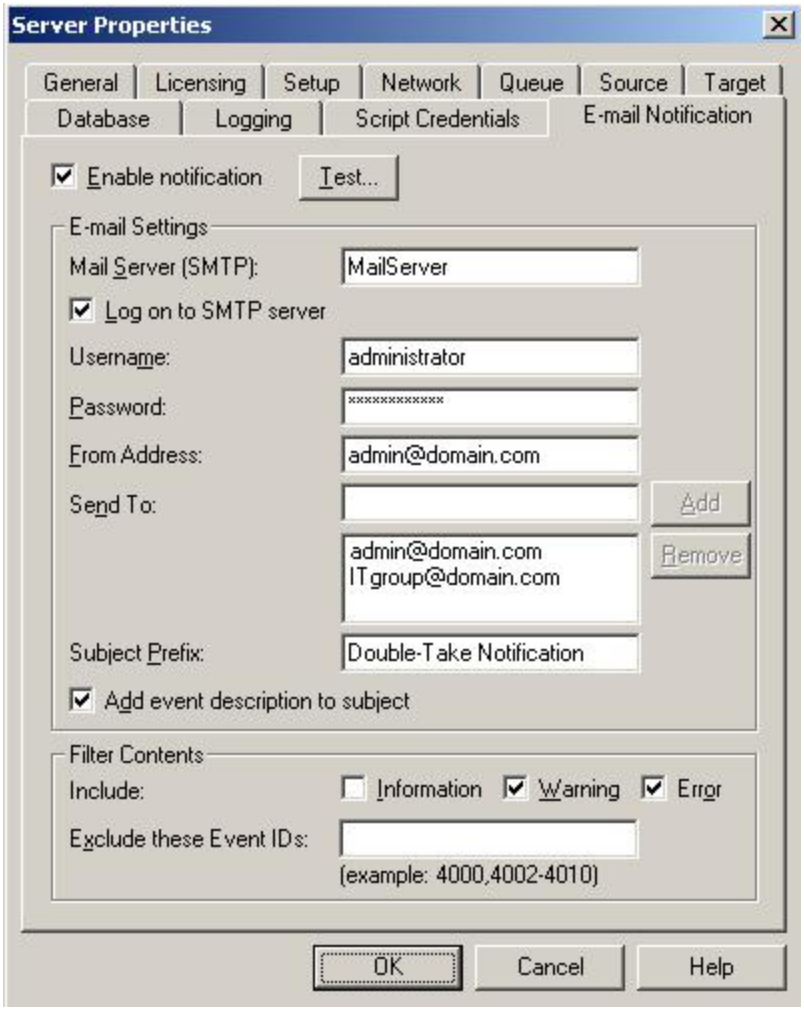

4. Select **Enable notification**.

Note: Any specified notification settings are retained when **Enable notification** is disabled.

5. Specify your e-mail settings.
   - **Mail Server (SMTP)**—Specify the name of your SMTP mail server.
   - **Log on to SMTP Server**—If your SMTP server requires authentication, enable **Log on to SMTP Server** and specify the **Username** and **Password** to be used for authentication. Your SMTP server must support the LOGIN authentication method to use this feature. If your server supports a different authentication method or does not support authentication, you may need to add the Double-Take Availability server as an authorized host for relaying e-mail messages. This option is not necessary if you are sending exclusively to e-mail addresses that the SMTP server is responsible for.
   - **From Address**—Specify the e-mail address that you want to appear in the From field of each Double-Take Availability e-mail message. The address is limited to 256 characters.
   - **Send To**—Specify the e-mail address that each Double-Take Availability e-mail message should be sent to and click **Add**. The e-mail address will be inserted into the list of addresses. Each address is limited to 256 characters. You can add up to 256 e-mail addresses. If you want to remove an address from the list, highlight the address and click **Remove**. You can also select multiple addresses to remove by Ctrl-clicking.
   - **Subject Prefix** and **Add event description to subject**—The subject of each e-mail notification will be in the format Subject Prefix : Server Name : Message Severity : Message ID : Message Description. The first and last components (Subject Prefix and Message Description) are optional. The subject line is limited to 150 characters.

     If desired, enter unique text for the **Subject Prefix** which will be inserted at the front of the subject line for each Double-Take Availability e-mail message. This will help distinguish Double-Take Availability messages from other messages. This field is optional.

     If desired, enable **Add event description** to subject to have the description of the message appended to the end of the subject line. This field is optional.
   - **Filter Contents**—Specify which messages that you want to be sent via e-mail. Specify **Information**, **Warning**, and/or **Error**. You can also specify which messages to exclude based on the message ID. Enter the message IDs as a comma or semicolon separated list. You can indicate ranges within

the list.

Note: You can test e-mail notification by specifying the options on the **E-mail Notification** tab and clicking **Test**. (By default, the test will be run from the machine where the Replication Console is running.) If desired, you can send the test message to a different e-mail address by selecting **Send To** and entering a comma or semicolon separated list of addresses. Modify the message text up to 1024 characters, if necessary. Click **Send** to test the e-mail notification. The results will be displayed in a message box. Click **OK** to close the message and click **Close** to return to the **E-mail Notification** tab.

E-mail notification will not function properly if the Event logs are full.

If an error occurs while sending an e-mail, a message will be generated. This message will not trigger an e-mail. Subsequent e-mail errors will not generate additional messages. When an e-mail is sent successfully, a message will then be generated. If another e-mail fails, one message will again be generated. This is a cyclical process where one message will be generated for each group of failed e-mail messages, one for each group of successful e-mail messages, one for the next group of failed messages, and so on.

If you start and then immediately stop the Double-Take service, you may not get e-mail notifications for the log entries that occur during startup.

By default, most virus scan software blocks unknown processes from sending traffic on port 25. You need to modify the blocking rule so that Double-Take Availability e-mail messages are not blocked.

---

6. Click **OK** to save the settings.

Statistics

Statistics logging is the process of taking snapshots of Double-Take Availability statistical data. The data can be written to a file for future use. Changes to the statistics file configuration are detected and applied immediately without restarting the Double-Take service.

The statistics log file created is a binary file. To view the log file, you must run the DTStat utility from the command prompt.

Sample DTStat output

```
================================
0/11/10 12:48:05:2040
================================
SYSTEMALLOCATOR::Total Bytes: 0
IQALLOCATOR::Total Bytes: 0
SECURITY::Logins : 1 FailedLogins : 0
KERNEL::SourceState: 2 TargetState: 1 Start Time: Tue Sep 11 12:45:26 2007
RepOpsGenerated: 436845 RepBytesGenerated: 0
MirOpsGenerated: 3316423 MirBytesGenerated: 108352749214952
   FailedMirrorCount: 0 FailedRepCount: 0
   ActFailCount: 0 TargetOpenHandles: 0 DriverQueuePercent: 0
TARGET:: PeerAddress: 10.10.1.104 LocalAddress: 10.10.1.104
   Ops Received: 25 Mirror Ops Received: 23
   Retries: 0 OpsDropped: 0 Ops Remaining: 0
   Orphan Files Removed: 0 Orphan Directories Removed: 0 Orphan Bytes
Removed: 0
   Bytes In Target Queue: 0 Bytes In Target Disk Queue: 0
   TasksSucceeded: 0 TasksFailed: 0 TasksIgnored: 0
SOURCE::autoDisConnects : 0 autoReConnects : 1
   lastFileTouched : /log/data_file
CONNECTION:: conPeerAddress: 10.10.1.104
   connectTime: Tue Sep 11 12:45:34 2007
   conState: 1 conOpsInCmdQueue: 0 conOpsInAckQueue: 0
   conOpsInRepQueue: 0 conOpsInMirQueue: 0 conBytesInRepQueue: 0
   conOpsTx: 27 conBytesInMirQueue: 0 conBytesTx: 14952687269
   conBytesCompressedTx: 14952
   conOpsRx: 201127 conBytesRx: 647062280 conResentOpCount: 0
conBytesInDiskQueue: 0
   conBandwidthLimit: 429496295 conBytesSkipped: 22867624
conMirrorBytesRemain: 0
   conMirrorPercent: 100.0%
   conTaskCmdsSubmitted: 0 conTaskCmdsQueued: 0
   conTasksSucceeded: 0 conTasksFailed: 0 conTasksIgnored: 0
```

Configuring the properties of the statistics file

1. Right-click a machine in the left pane of the Replication Console and select **Properties**.
2. Select the **Logging** tab.

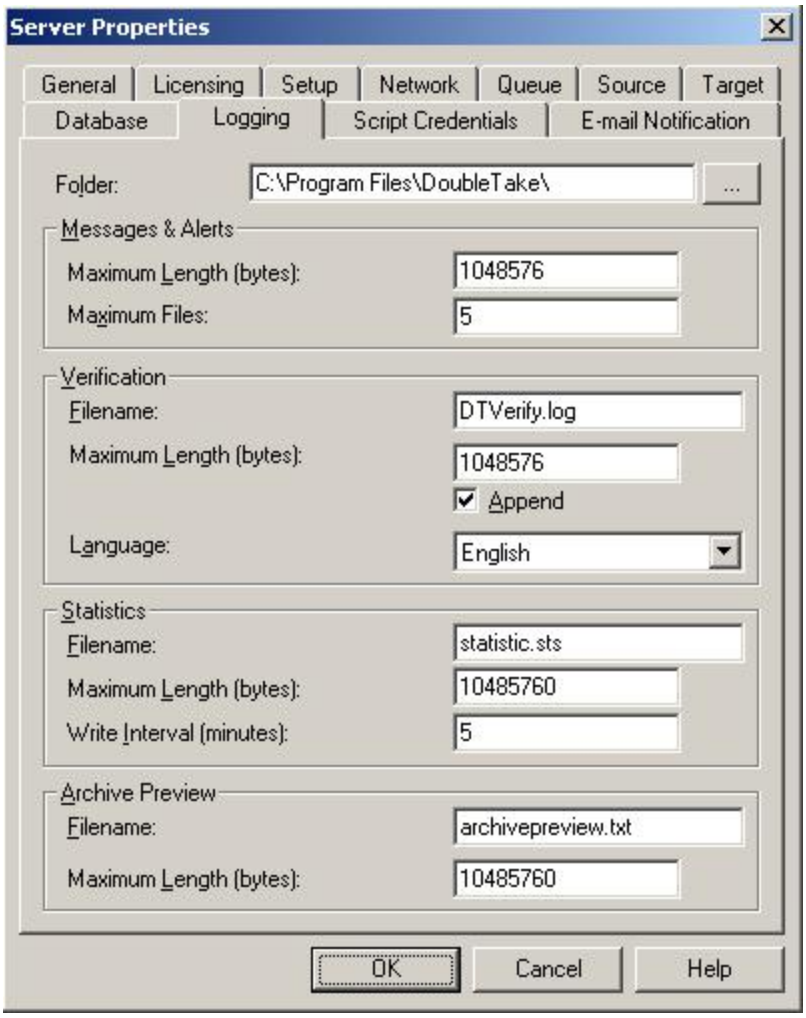

3. At the top of the tab, specify the **Folder** where the log files for messages, alerts, verification, and statistics will be saved.
4. Under **Statistics**, specify the following information.
   - **Filename**—The name of the statistics log file. The default file name is statistic.sts.
   - **Maximum Length**—The maximum length of the statistics log file. The default maximum length is 10 MB. Once this maximum has been reached, Double-Take Availability begins overwriting the oldest data in the file.

- **Write Interval**—The frequency in which Double-Take Availability writes the statistical data to the statistics log file. The default is every 5 minutes.

5. Select the **Setup** tab.
6. Verify that **Log Statistics Automatically** is enabled. If disabled, statistics will not be logged.
7. Click **OK** to save the settings.

Viewing the statistics file

The statistics log file created is a binary file. To view the log file, you must run the DTStat utility from a command prompt. From the directory where Double-Take Availability is installed, run the DTStat command.

Command

DTSTAT

Description

Starts the DTStats statistics logging utility from a command prompt

Syntax

DTSTAT [-p][-i <*interval*>][-t <*filename*>] [-f <*filename*>] [-s <*filename*>] [-st <*filename*>][-IP <*address*>] [-START <*mm*/*dd*/*yyyy hh*:*mm*>][-STOP <*mm*/*dd*/*yyyy hh*:*mm*>] [-SERVER <*ip_address*> <*port_number*>]

Options

- -p—Do not print the output to the screen
- -i *interval*—Refresh from shared memory every interval seconds
- -t *filename*—Save the data from memory to the specified binary file filename
- -f *filename*—Reads from a previously saved binary file, filename, that was generated using the -t option instead of reading from memory
- -s *filename*—Saves only the connection data from the data in memory to an ASCII, comma-delimited file, filename
- -st *filename*—Saves only the target data from the data in memory to an ASCII, comma-delimited file, filename
- -f *filename1* -s *filename2*—Saves only the connection data from a previously saved binary file, filename1, to an ASCII, comma-delimited file, filename2
- -f *filename1* -st *filename2*—Saves only the target data from a previously saved binary file, filename1, to an ASCII, comma-delimited file, filename2
- -*IP address*—Filters out the specified address in the IP address field and prints only those entries. Specify more than one IP address by separating them by a comma.
- -START *mm*/*dd*/*yyyy hh*:*mm*—Filters out any data prior to the specified date and time

- -STOP ***mm*/*dd*/*yyyy hh*:*mm***—Filters out any data after the specified date and time
- -SERVER ***ip_address port_number***—Connects DTStat to the specified IP address using the specified port number instead of to the local machine

Examples

- DTStat -i 300
- DTStat -p -i 300 -t AlphaStats.sts
- DTStat -f AlphaStats.sts -s AlphaStats.csv -start 02/02/2007 09:25
- DTStat -server 206.31.4.51 1106

Notes

- This command is not case-sensitive.
- If no options are specified, DTStat will print the output to the screen at an interval of every one second.
- If the statistics are not changing, DTStat will discontinue writing until statistics begin updating again.

Statistics

The following table identifies the Double-Take Availability statistics.

Note: The categories you see will depend on the function of your server (source, target, or both).

If you have multiple IP addresses connected to one target server, you will see multiple Target sections for each IP address.

If you convert your statistics output to an ASCII, comma-delimited file using the dtstat -s option, keep in mind the following differences.

- The statistic labels will be slightly different in the ASCII file than in the following table.
- The statistics will appear in a different order in the ASCII file than in the following table.
- The statistics in the Target Category in the following table are not included in the ASCII file.
- The Kernel statistic Target Open Handles is not included in the ASCII file.
- The ASCII file contains a Managed Pagefile Alloc statistic which is no longer used.

Date/Time Stamp

The date and time that the snapshot was taken. This is the date and time that each statistic was logged. By default, these are generated once a second, as long as there are statistics being generated. If mirroring/replication is idle, then DTStat will be idle as well.

System Allocator, Total Bytes

The number of bytes currently allocated to the system pagefile

IQAllocator, Total Bytes

The number of bytes currently allocated to the intermediate queue

Security, Logins

The number of successful login attempts

Security, Failed Logins

The number of failed login attempts

Kernel, SourceState

- 0—Source is not running
- 1—Source is running without the replication driver
- 2—Source is running with the replication driver

Kernel, TargetState

- 0—Target is not running
- 1—Target is running

Kernel, Start Time

Date and time stamp indicating when the Double-Take service was loaded

Kernel, RepOpsGenerated

The number of replication operations generated by the file system driver. An op is a file system operation. Double-Take Availability replicates data by sending the file system operations across the network to the target. RepOpsGenerated indicates the number of file system operations that have been generated by replication.

Kernel, RepBytesGenerated

The number of replication bytes generated by the file system driver. This is the number of bytes generated during replication. In other words, this is roughly the amount of traffic being sent across the network that is generated by replication. It does not take into account TCP/IP overhead (headers and such).

Kernel, MirOpsGenerated

The number of mirror operations transmitted to the target. Mirroring is completed by transmitting the file system operations necessary to generate the files on the target. This statistic indicates the number of file system operations that were transmitted during the initial mirror. It will continue to increase until the mirror is complete. Any subsequent remirrors will reset this field to zero and increment from there.

Kernel, MirBytesGenerated

The number of mirror bytes transmitted to the target. This is the number of bytes generated during mirroring. In other words, this is roughly the amount of traffic being sent across the network that is generated by the mirror. It does not take into account TCP/IP overhead (headers and such). Again, any subsequent remirror will reset this field to zero and increment from there.

Kernel, FailedMirrorCount

The number of mirror operations that failed due to an error reading the file from the disk

Kernel, FailedRepCount

The number of replication operations that failed due to an error reading the file from the disk

Kernel, ActFailCount

The number of activation code failures when loading the source or target. Activation codes can be bad for reasons such as: expiration of evaluation codes, duplicate codes, incorrect codes, etc.

Kernel, TargetOpenHandles

The number of handles currently open on the target

Kernel, DriverQueuePercent

The amount of throttling calculated as a percentage of the stop replicating limit

Target, PeerAddress

The IP address of the source machine

Target, LocalAddress

The IP address of the target machine.

Target, Ops Received

The total number of operations received by this machine as a target since the Double-Take service was loaded

Target, Mirror Ops Received

The total number of mirror operations received by this machine as a target since the Double-Take service was loaded. This number does not reset to zero for remirrors.

Target, Retries

The number of retries performed before all operations were completed

Target, OpsDropped

The number of operations skipped during a difference mirror. During a difference mirror, if Double-Take Availability detects that there have been no changes to a file, then it will indicate the number of operations it did not send for this file in this field.

Target, Ops Remaining

The total number of operations that are left in the target queue

Target, Orphan Files Removed

The number of orphan files removed from the target machine

Target, Orphan Directories Removed

The number of orphan directories removed from the target machine

Target, Orphan Bytes Removed

The number of orphan bytes removed from the target machine

Target, Bytes In Target Queue

The number of bytes currently in the system memory queue on the target

Target. Bytes In Target Disk Queue

The number of bytes currently in the disk queue on the target

Target, TasksSucceeded

The number of task commands that have succeeded on the target

Target, TasksFailed

The number of task commands that have failed on the target

Target, TasksIgnored

The number of task commands that have been ignored on the target

Source, autoDisConnects

The number of automatic disconnects since starting Double-Take Availability. Auto-disconnects occur because the source no longer sees the target This could be because the connection between the two has failed at some point or because the target machine data is changing on the source faster than the source can get the data to the target. This field tracks the number of times an auto-disconnect has occurred since the Double-Take service was started.

Source, autoReConnects

The number of automatic reconnects since starting Double-Take Availability. Auto-reconnect occurs after a target machine is back online. This field tracks the number of times an auto-reconnect has happened since the Double-Take service was started.

Source, lastFileTouched

The last filename that had a replication operation executed

Connection, conPeerAddress

The IP address of the target machine

Connection, connectTime

The time that this connection was established

Connection, conState

The state of the active connection

- 0—None. This indicates a connection has not been established. Statistics are still available for the source and target machines.
- 1—Active. This indicates that the connection is functioning normally and has no scheduling restrictions imposed on it at this time. (There may be restrictions, but it is currently in a state that allows it to transmit.)
- 2—Paused. This indicates a connection that has been paused.
- 4—Scheduled. This indicates a connection that is not currently transmitting due to scheduling restrictions (bandwidth limitations, time frame limitations, and so on).
- 8—Error. This indicates a connection that is not transmitting because something has gone wrong (for example, lost connection).

Only the Scheduled and Error states can coexist. All other states are mutually exclusive. Statistics will display a conState of 12 when the connection is in both a scheduled and an error state because this is the sum of the two values (4 + 8).

Connection, conOpsInCmdQueue

The number of operations waiting to be executed on the target

Connection, conOpsInAckQueue

The number of operations waiting in the acknowledgement queue. Each operation that is generated receives an acknowledgement from the target after that operation has been received by the target. This statistic indicates the number of operations that have yet to receive acknowledgement of receipt.

Connection, conOpsInRepQueue

The number of replication operations currently waiting to be executed on the target

Connection, conOpsInMirQueue

The number of mirror operations currently waiting to be executed on the target

Connection, conBytesInRepQueue

The number of replication bytes remaining to be transmitted to the target

Connection, conOpsTx

The number of operations transmitted to the target. This is the total number of operations that Double-Take Availability has transmitted as a source. In other words, the cumulative number of operations transmitted by this source to all connected targets.

Connection, conBytesInMirQueue

The number of mirror bytes remaining to be transmitted to the target

Connection, conBytesTx

The number of bytes transmitted to the target. This is the total number of bytes that Double-Take Availability has transmitted as a source. In other words, the cumulative number of bytes transmitted by this source to all connected targets.

Connection, conBytesCompressedTx

The number of compressed bytes transmitted to the target.

Connection, conOpsRx

The number of operations received by the target. The number of operations that the target for this connection (as indicated by the IP address field) has received from this source.

Connection, conBytesRx

The number of bytes received by the target. The number of bytes that the target for this connection (as indicated by the IP address field) has received from this source.

Connection, conResentOpCount

The number of operations resent because they were not acknowledged

Connection, conBytesInDiskQueue

The number of bytes in the source disk queue

Connection, conBandwidthLimit

The amount of bandwidth that may be used to transfer data

Connection, conBytesSkipped

The number of bytes skipped during a difference mirror. During a difference mirror, if Double-Take Availability detects that there have been

no changes to a file, then it will indicate the number of bytes it did not send for this file in this field.

Connection, conMirrorBytesRemaining

The number of mirror bytes remaining to be transmitted

Connection, conMirrorPercent

The percentage of the mirror that has been completed. This field is determined if the replication set size was calculated.

Connection, conTaskCmdsSubmitted

The number of task commands that have been submitted on the source

Connection, conTaskCmdsQueued

The number of task commands that have been queued on the source

Connection, conTasksSucceeded

The number of task commands that have succeeded on the source

Connection, conTasksFailed

The number of task commands that have failed on the source

Connection, conTasksIgnored

The number of task commands that have been ignored on the source

Performance Monitor

Performance Monitor is the Windows graphical tool for measuring performance. It provides charting, alerting, and reporting capabilities that reflect both current activity and ongoing logging. Double-Take Availabilitystatistics are available through the Performance Monitor.

- Monitoring Performance Monitor statistics
- Performance Monitor statistics

Monitoring Performance Monitor statistics

1. To access the Performance Monitor, select **Start**, **Programs**, **Administrative Tools**, **Performance**.
2. Specify the data to monitor by right-clicking and selecting **Add** or using the **Add** button on the toolbar.
3. Choose one of the following Double-Take Availability Performance Objects.
    - Double-Take Connection
    - Double-Take Kernel
    - Double-Take Security
    - Double-Take Source
    - Double-Take Target
4. Select the statistics you want to monitor, and click **Add**.

For additional information and details on the Performance Monitor, see your Windows reference guide.

Performance Monitor statistics

The following table identifies the Double-Take Availability Performance Monitor statistics.

Note: If you have multiple IP addresses connected to one target server, you will see multiple Target statistic sections for each IP address.

Connection, Bandwidth Limit

The amount of bandwidth that may be used to transfer data

Connection, Bytes in disk queue

The number of bytes in the source disk queue

Connection, Bytes in replication queue

The number of replication bytes in the source queue

Connection, Bytes in mirror queue

The number of mirror bytes in the source queue

Connection, Bytes received

The number of bytes received by the target since the last Performance Monitor refresh

Connection, Bytes transferred

The number of bytes transmitted from the source

Connection, Compressed bytes transferred

The number of compressed bytes transmitted from the source

Connection, Operations in acknowledgement queue

The number of operations waiting in the source acknowledgement queue

Connection, Operations in command queue

The number of operations waiting in the source command queue

Connection, Operations in mirror queue

The number of mirror operations in the source queue

Connection, Operations in replication queue

The number of replication operations in the source queue

Connection, Operations received

The number of operations received by the target since the last Performance Monitor refresh

Connection, Operations resent

The number of operations re-sent since the last time the Double-Take service was restarted on the source

Connection, Operations transmitted

The number of operations transmitted from the source

Connection, Task commands queued

The number of task commands queued on the source

Connection, Task commands submitted

The number of task commands submitted on the source

Connection, Tasks failed

The number of task commands that have failed to execute on the source

Connection, Tasks ignored

The number of task commands that have been ignored on the source

Connection, Tasks succeeded

The number of task commands that have succeeded on the source

Kernel, Activation code failures

The number of activation code failures when loading the source or target, since the last time the Double-Take service was restarted on the source

Kernel, Double-Take queue memory usage

The amount of system memory in use by the Double-Take Availability queue

Kernel, Driver Queue Percent

The amount of throttling calculated as a percentage of the stop replicating limit

Kernel, Failed mirror operations

The number of mirror operations on the source that failed due to an error reading the file from the disk

Kernel, Failed replication operations

The number of replication operations on the source that failed due to an error reading the file from the disk

Kernel, Mirror Kbytes generated

The number of mirror kilobytes transmitted from the source

Kernel, Mirror operations generated

The number of mirror operations transmitted from the source

Kernel, Open Target Handles

The number of handles currently open on the target.

Kernel, Replication Kbytes generated

The number of replication kilobytes generated on the source by the file system driver

Kernel, Replication operations generated

The number of replication operations generated on the source by the file system driver

Security, Failed logins

Number of failed login attempts since the last time the Double-Take service was restarted

Security, Successful logins

Number of successful login attempts since the last time the Double-Take service was restarted

Source, Auto disconnects

The number of automatic disconnects since the last time the Double-Take service was restarted on the source

Source, Auto reconnects

The number of automatic reconnects since the last time the Double-Take service was restarted on the source

Target, Bytes in Disk Queue

The number of bytes in the target disk queue

Target, Bytes in Queue

The number of bytes in the system memory and disk queues

Target, Mirror operations received

The number of mirror operations received on the target

Target, Operations received

The number of operations received on the target

Target, Ops Dropped

The number of operations dropped on the target since the last time the Double-Take service was restarted on the target

Target, Ops Remaining

The number of operations on the target remaining to be applied

Target, Orphan Bytes

The number of orphan bytes removed from the target

Target, Orphan Directories

The number of orphan directories removed from the target

Target, Orphan Files

The number of orphan files removed from the target

Target, Retries

The number of retries performed on the target since the last time the Double-Take service was restarted on the target

Target, Tasks failed

The number of task commands that have failed on the target.

Target, Tasks ignored

The number of task commands that have been ignored on the target

Target, Tasks succeeded

The number of task commands that have succeeded on the target

SNMP

SNMP, Simple Network Management Protocol, is the Internet's standard for remote monitoring and management of hosts, routers and other nodes and devices on a network. Double-Take Availability provides an SNMP sub-agent that can be managed from an SNMP Management Console.

Double-Take Availability installs two components to work with SNMP.

- The sub-agent is a program that installs and runs on the same machine as Double-Take Availability and gathers statistics, data, and traps. The sub-agent forwards the information to the SNMP agent, which relays the information to the manager. The Double-Take Availability SNMP sub-agent is included in the Double-Take Availability installation program.
- A Double-Take Availability MIB file is placed on the administrator's machine so that the Management Console can interpret the data sent from the sub-agent. The Double-Take Availability .mib file is dt.mib and meets SNMP standards.

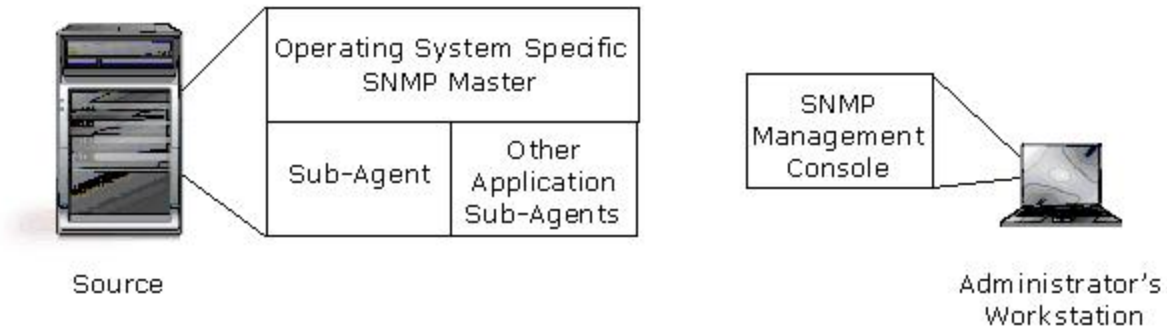

- Configuring SNMP on your server
- SNMP traps
- SNMP statistics

Configuring SNMP on your server

SNMP must be installed on a server before Double-Take Availability in order for the Double-Take AvailabilitySNMP components to be added during the Double-Take Availabilityinstallation. If SNMP is installed on a server after Double-Take Availability is installed, run a repair install to install the SNMP components.

The Double-Take Availability .mib file will need to be loaded into your SNMP Management Console. Depending on the type of console you are using, this process might include compiling the .mib file. Reference your SNMP Management Console documentation for additional information.

SNMP traps

The following table lists the Double-Take Availability SNMP traps.

Kernel, dttrapKernelStarted

Double-Take Availability has started

Kernel, dttrapKernelStopped

Double-Take Availability has stopped

License, dttrapLicenseViolationStartingSource

The source cannot be started due to a license violation

License, dttrapLicenseViolationOnNetwork

A Double-Take Availability serial number conflict was identified on the network

Source, dttrapSourceStarted

Double-Take Availability source component has started

Source, dttrapSourceStopped

Double-Take Availability source component has stopped

Target, dttrapTargetStarted

Double-Take Availability target component has started

Target, dttrapTargetStopped

Double-Take Availability target component has stopped

Connection, dttrapConnectionRequested

The source has requested a connection to the target

Connection, dttrapConnectionRequestReceived

The target has received a connection request from the source

Connection, dttrapConnectionSucceeded

The source to target connection has been established

Connection, dttrapConnectionPause

The source to target connection has paused

Connection, dttrapConnectionResume

The source to target connection has resumed

Connection, dttrapConnectionFailed

The source to target connection was not successful

Connection, dttrapConnectionLost

The source to target connection has been disconnected

Connection, dttrapMemoryLimitReached

The Double-Take Availability memory pool limit has been reached

Connection, dttrapMemoryLimitRemedied

The memory pool usage is below the maximum limit specified

Connection, dttrapAutoReconnect

Auto-reconnect needs to make a new connection

Connection, dttrapScheduledConnectStart

A scheduled connection has been established

Connection, dttrapScheduledConnectEnd

A scheduled end connection has been reached and the connection has been disconnected

Connection, dttrapAutoDisconnectWriteQueue

Auto-disconnect has forced the queue to be written to disk

Connection, dttrapAutoDisconnectPauseTransmission

Auto-disconnect requested that the source pause any operation (create, modify, or delete) sending

Connection, dttrapAutoDisconnectEndConnection

Auto-disconnect has intentionally dropped the connection

Connection, dttrapAutoDisconnectShutdown

Auto-disconnect forced Double-Take Availability to shutdown

Replication, dttrapReplicationStart

Replication has started

Replication, dttrapReplicationStop

Replication has stopped

Mirroring, dttrapMirrorStart

Mirroring has started

Mirroring, dttrapMirrorStop

Mirroring has stopped

Mirroring, dttrapMirrorPause

Mirroring has paused

Mirroring, dttrapMirrorResume

Mirroring has resumed

Mirroring, dttrapMirrorEnd

Mirroring has ended

Verification, dttrapVerificationStart

Verification has started

Verification, dttrapVerificationEnd

Verification has ended

Verification, dttrapVerificationFailure

Verification has failed

Restoration, dttrapRestoreStarted

Restoration has started

Restoration, dttrapRestoreComplete

Restoration is complete

Replication Sets, dttrapRepSetModified

Replication has been modified

Failover, dttrapFailoverConditionMet

Manual intervention is required because failover has detected a failed source machine

Failover, dttrapFailoverInProgress

Failover is occurring

Failover, dttrapTargetFull

The target is full

SNMP statistics

The following table lists the Double-Take Availability SNMP statistics.

General, dtUpTime

Time in seconds since Double-Take Availability was last started

General, dtCurrentMemoryUsage

Amount of memory allocated from the Double-Take Availability memory pool

General, dtMirOpsGenerated

The number of mirror operations (create, modify, or delete) that have been transmitted by the mirroring process

General, dtMirBytesGenerated

The number of bytes that have been transmitted by the mirroring process

General, dtRepOpsGenerated

The number of operations (create, modify, or delete) that have been transmitted by the replication process

General, dtRepBytesGenerated

The number of bytes that have been transmitted by the replication process

General, dtFailedMirrorCount

The number of operations that failed to mirror because they could not be read on the source

General, dtFailedRepCount

The number of operations that failed to be replicated because they could not be read on the source

General, dtActFailCount

The number of activation code errors

General, dtAutoDisCount

The number of auto-disconnects

General, dtAutoReCount

The number of auto-reconnects

General, dtDriverQueuePercent

The amount of throttling calculated as a percentage of the stop replicating limit

Source, dtSourceState

0—Source is not running

1—Source is running without the replication driver

2—Source is running with the replication driver.

Target, dtTargetState

0—Target is not running

1—Target is running

Target, dtRetryCount

The number of file operations that have been retried

Target, dtOpsDroppedCount

The number of file operations that have failed and will not be retried

Security, dtLoginCount

The number of successful logins

Security, dtFailedLoginCount

The number of unsuccessful logins

Connection, dtConnectionCount

The number of active connections between machines

Connection, dtconIpAddress

The IP address of the connected machine. If at the source, then the IP address of the target. If at the target, then the IP address of the source.

Connection, dtconConnectTime

The duration of time since the connection was first established

Connection, dtconState

The state of the active connection

0—None. This indicates a connection has not been established. Statistics are still available for the source and target machines.

1—Active. This indicates that the connection is functioning normally and has no scheduling restrictions imposed on it at this time. (There may be restrictions, but it is currently in a state that allows it to transmit.)

2—Paused. This indicates a connection that has been paused.

4—Scheduled. This indicates a connection that is not currently transmitting due to scheduling restrictions (bandwidth limitations, time frame limitations, and so on).

8—Error. This indicates a connection that is not transmitting because something has gone wrong (for example, lost connection).

Only the Scheduled and Error states can coexist. All other states are mutually exclusive. SNMP will display a dtconState of 12 when the connection is in both a scheduled and an error state because this is the sum of the two values (4 + 8).

Connection, dtconOpsInCmdQueue

The number of operations (create, modify, or delete) in the retransmit queue on the source

Connection, dtconOpsInAckQueue

The number of operations (create, modify, or delete) waiting for verification acknowledgements from the target

Connection, dtconOpsInRepQueue

The number of replication operations (create, modify, or delete) in the queue

Connection, dtconOpsInMirQueue

The number of mirror operations (create, modify, or delete) in the queue

Connection, dtconBytesInRepQueue

The number of bytes in the replication queue

Connection, dtconBytesInMirQueue

The number of bytes in the mirror queue

Connection, dtconOpsTx

The total number of operations (create, modify, or delete) transmitted to the target

Connection, dtconBytesTx

The total number of bytes transmitted to the target

Connection, dtconBytesCompressedTx

The total number of compressed bytes transmitted to the target

Connection, dtconOpsRx

The total number of operations (create, modify, or delete) received from the target

Connection, dtconBytesRx

The total number of bytes received from the target

Connection, dtconResentOpCount

The number of operations that were resent because of acknowledgement errors

---

Error codes

The following table contains error codes that you may see in the various user interfaces or in log files.

-1 Unknown error code (generated when a command failed but the failure is not linked to a pre-defined error code)

-101 Invalid parameter was supplied

-102 Command is not a valid or the syntax is incorrect

-103 Double-Take Availability source module is not loaded

-104 No Double-Take Availability source identified

-105 Double-Take Availability target module is not loaded

-106 Connection already established

-107 Connection does not exist

-108 Mirror currently active

-109 Server does not exist or could not be located

-110 Server is not responding

-111 Double-Take Availability is running

-112 Unknown connection error

-113 Mirror already active

-114 Date is invalid - valid format is mm/dd/yy

-115 Time is invalid - valid format is hh:mm

-116 Invalid option supplied

-117 Mirror is not paused

-118 Connection is not paused

-119 Connection does not exist

-120 Connection already connected

-121 Mirror is not running

-122 Replication set exists

-123 Replication set does not exist

-124 No replication set has been selected

-125 Connection is replicating

-126 Connection is not replicating

-127 Replication set is enabled

-128 Schedule is not defined

-129 Replication set is changed

-130 Replication set is in use

-131 No Double-Take Availability target identified

-132 Memory is low

-133 Memory is sufficient

-134 Replication is pending

-135 Invalid option supplied

-136 Replication set rule does not exist

-137 Mirror queue is full

-138 Insufficient security access

-139 Schedule command is invalid

-140 Source path is invalid

-141 Replication set is not changed

-142 Insufficient source security access

-143 Invalid statistics file

-144 Replication set not saved

-145 Connection failed

-146 Cleaner option is not enabled

-147 Target mirror capacity high threshold is met

-148 Target mirror capacity low threshold is met

-149 New option applied

-150 Target is restarted

-151 Replication is out of memory

-152 Write access is blocked on the volume

-153 This error code could be one of two errors. 1) Compression level is not supported, or server does not support compression 2) Transmission is paused

-154 Transmission is active

-155 Target does not support the command

-156 Command conversion to accommodate a different Double-Take Availability version has failed

-157 Incompatible source and target Double-Take Availability versions

-158 Incompatible source and target operating system versions

-159 NAS server to non-NAS server is not a supported configuration

-160 Target module is not loaded

-161 Operation or command is not supported

-162 Target is paused

-163 Target is pending

-164 Target is active

-165 Target is retrying operations

-166 Target is no longer retrying operations

-167 Restore required state is unknown

-168 Not a valid failover source

-169 Failover login failed

-170 Feature is not supported

-171 Command is not supported

-172 Target queue log file error

-173 Target disk is full

-174 Target disk has sufficient disk space

-175 Error reading from or writing to the queue log file

-176 Memory-based queue is in use

-177 Disk-based queue is in use

-178 Restore is required

-179 ID the driver supplied to the service is invalid

-180 Child path is blocked

-181 Parent path is blocked

-182 Target path blocking is disabled

-183 Connection ID specified is invalid

-184 No command objects are in the queue

-185 Target is discarding operations from the target queue

-186 Target is not discarding operations from the target queue

-187 Schedule is paused

-188 Schedule is resumed

-189 Target state has changed

-190 Target name has changed

-201 Monitor name exists

-202 Monitor name does not exist

-203 Monitor configuration exists

-204 Monitor configuration does not exist

-205 Monitor configuration is in use

-206 Monitor configuration is not in use

-207 Source is online

-208 Source is offline

-209 Server is not failed over

-210 Server is failed over

-211 Server is not being monitored

-212 Failback is in progress

-213 IP address placeholders on the target are unavailable

-214 Target NIC was not found

-215 Source module is not loaded

-216 Failed to set the source state

-217 Unable to ping source

-218 Invalid argument

-219 Recovery is busy

-220 Invalid command

-221 Recovery is started

-222 Script failed to start

-223 Script timeout met

-224 No replication timeout met - connection is bad

-225 Invalid path

-226 Kernel module is not loaded

-2201 Error communicating with e-mail server

-2202 Error connecting to e-mail server

-2203 E-mail notification is disabled

-2204 E-mail notification is enabled

-2205 E-mail notification requires Internet Explorer version 5.0 and WMI

-2206 E-mail notification requires Internet Explorer version 5.0 (E-mail notification no longer requires Internet Explorer 5.0 or later. If you receive this error, contact technical support.)

-2207 Error sending e-mail

-2208 Error sending test e-mail

-2209 WMI error connecting to e-mail server

-2210 E-mail notification requires WMI

-2211 Event Viewer settings for e-mail notification are invalid

-2212 E-mail notification setting is invalid

-2213 E-mail notification address exists

-2214 E-mail notification alert ID is invalid

-2215 E-mail notification format is invalid

-2216 E-mail notification address does not exist

-2217 E-mail notification address notification list is empty

-2218 E-mail warning is not set

-2219 E-mail test warning is not set

-2200 E-mail notification is functioning properly

-2301 Bandwidth limiting time exists

-2302 Bandwidth limiting name exists

-2303 Bandwidth limit not found

-2304 Bandwidth limit day is invalid

-2305 Bandwidth limit label is invalid

-2401 Snapshot module is not loaded

-2402 Error reading the snapshot .dll

-2403 Snapshot not found

-2404 No snapshot connections found

-2501 Full-server functionality is disabled

-2502 No full-server interface available

Failover

Your failover process will depend on the type of workload you are protecting.

- Failing over data workloads, application workloads configured for identity failover, and cluster workloads
- Full-server workload failover
- Failing over application workloads configured for DNS failover
- Virtual workload failover

Failing over data workloads, application workloads configured for identity failover, and cluster workloads

The failover process, including script processing, can be tested at any time. To force unavailability, disconnect the network cable from a monitored machine, wait for the **Time to Fail** counter to decrease to zero and failover begins. To avoid the countdown delay, highlight the monitored machine name in the Failover Control Center window and select **Failover**.

If **Manual Intervention** is enabled, the Failover Control Center will prompt you when a failure occurs. Click **Cancel** to abort the failover process. (If necessary, you can initiate failover later from the Failover Control Center.) Click **OK** to proceed with failover.

Note: If the Failover Control Center is not running at the time the failure occurs, the manual intervention dialog box will appear the next time the Failover Control Center is started.

When a failure occurs, an alert is forwarded to the Windows event log. You can then start the Failover Control Center and respond to the manual intervention prompt.

If SNMP is installed and configured, an SNMP trap is also generated. When using a third-party SNMP manager, an e-mail or page can be generated to notify you of the failure.

You will be prompted to specify what data you want to use on the target.

Select an option based on the table below. You may want to check the amount of data in queue on the target by reviewing the Statistics or Performance Monitor.

Apply Data in Target Queues Then Failover

All of the data in the target queue will be applied before failover begins.

Advantages—All of the data that the target has received will be applied before failover begins.

Disadvantages—Depending on the amount of data in queue, the amount of time to apply all of the data could be lengthy.

Discard Data in Target Queues and Failover Immediately

All of the data in the target queue will be discarded and failover will begin immediately.

Advantages—Failover will occur immediately.

Disadvantages—Any data in the target queue will be lost.

Revert to Last Good Snapshot if Target Data State is Bad

If the target data is in a bad Double-Take Availability state, Double-Take Availability will automatically revert to the last good Double-Take Availability snapshot before failover begins. If the target data is in a good state, Double-Take Availability will not revert the target data. Instead, Double-Take Availability will apply the data in the target queue and then failover.

Advantages—Good data on the target is guaranteed to be used.

Disadvantages—If the target data state is bad, you will lose any data between the last good snapshot and the failure.

After failover is complete, clients will be rerouted to the target, which is standing in for the source.

Exchange Note: Users using Outlook or Outlook Web Access to receive e-mail can connect after the changes have propagated through your environment. Users that had Outlook open during failover will need to restart the Outlook client (excluding Outlook Web Access clients on a LAN). Additionally, those users using Outlook Web Access or Outlook 2007 may see a security alert because the security certificate has the source server name but Exchange is now on the target. Click **Allow** or **OK** to dismiss the alert.

You will not be able to log in to the domain from the source Exchange server after failover because the target has assumed the source server's host Service Principal Name so that Outlook Web Access can use the source name. If you need to log in to the domain and Outlook Web Access is not needed, contact technical support for a workaround.

If you SMTP gateway is configured to send e-mail to a specific IP address that address is not failed over to the target, you will need to update the IP address after failover.

Mail stores or storage groups created after a failover will not be failed back.

SQL Note: After failover, linked databases in the SQL instance will be unavailable until the service master key is updated. You will need to run the command "alter service master key force regenerate" against the SQL target server to reset the service master key and then remove and re-add the linked servers into the target SQL instance.

After failover with a snapshot of a SQL database-only server, the SQL services on the target server are stopped and the databases are not mounted. You will need to manually start the MSSQLServer service for each instance on the target server and then manually attach the databases.

After failing over SQL 2008, Rich Internet Applications created using ADO.net 2 may not connect.

After failing over SQL 2008, you may not be able to take the SQL database offline. If this occurs, stop and restart the SQL Server Management Studio application, and then you can take the database offline.

After failing over in a SQL workgroup, you will not be able to connect to the source server instance of SQL. You can work around this issue by creating an alias on the target with the source server’s name.

BlackBerry Note: There is a potential to lose BlackBerry instant messages and instant message contacts when an Exchange mailbox is moved, depending on the type of BlackBerry hand-held device. You should back up

instant messages through the BlackBerry desktop software or have the BlackBerry Enterprise server configured to audit all instant messages through a policy.

Full-server workload failover

Full-server failover can be initiated through the Full-Server Failover Manager client or by using a command line interface.

- Failing over using the Full-Server Failover Manager
- Failing over from the command line

Failing over using the Full-Server Failover Manager

When a failover condition is met, you will want to start failover. Additionally, you can start it without a failover condition, as long as protection is enabled. For example, you may want to force failover when upgrading to a better source server.

1. To start failover, click **Failover**.
2. If Double-Take Availability determines there is a possibility that the data on the target is incomplete, you will be warned before failover begins. If you proceed with failover, the state of the source will be unknown until failover is complete. The best case scenario would be a missing data file, while the worst case scenario would be missing system state data that causes the server to be unusable or unbootable. Select your failover options.

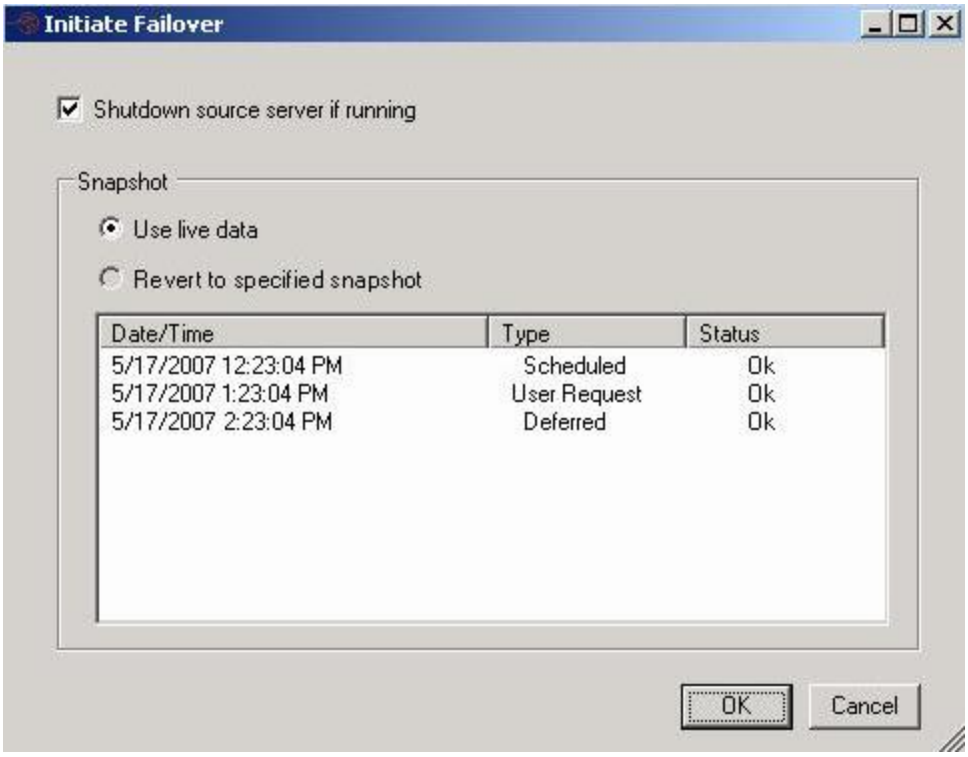

- **Shutdown source server if running**—If the source is still running, Full-Server Failover Manager can stop it. Although, if Full-Server Failover Manager cannot communicate with the source, the shutdown command will fail. This option prevents network conflicts in those cases where the source and target are still both running and communicating, such as a forced failover.
- **Use live data**—Select this option to use the data on the target at the time of failover.

- **Revert to specified snapshot**—Select this option, and then select a snapshot. The data on the target will be reverted to the selected snapshot. This option will not be available if there are no snapshots on the target or if the target does not support snapshots. To help you understand what snapshots are available, use the **Type** and **Status** columns. The **Status** indicates the state of the connection between the source and target at the time the snapshot was taken. The **Type** indicates the kind of snapshot.
  - **Scheduled**—This snapshot was taken as part of a periodic snapshot.
  - **Deferred**—This snapshot was taken as part of a periodic snapshot, although it did not occur at the specified interval because the connection between the source and target was not in a good state.
  - **User Request**—This snapshot was taken manually by a user.

3. Click **OK** to initiate failover. Monitor the failover percentage as shown in the Protection Status. At the end of failover, the target will be rebooted automatically. After the reboot, the target will no longer exist, since it will become the source.

Note: If you are failing over a cluster node, it is possible that volumes may lose their drive letter assignments. If a clustered application fails to start after failover and the disk signature has changed, check the drive letter assignments under the Disk Management utility and re-create drive letter assignments as needed.

Because the Windows product activation is dependent on hardware, you may need to reactivate your Windows registration after failover. In most cases when you are using Windows 2003, you can follow the on-screen prompts to complete the reactivation. However, when you are using Windows 2008, the reactivation depends on your licensing type. If a Windows 2008 target comes online after failover with an activation failure, use the steps appropriate for your license type.

- **Retail licensing**—Retail licensing allows the activation of a single operating system installation.
  1. Open the **System** applet in Windows **Control Panel**.
  2. Under **Windows activation** at the bottom of the page, click **Change product key**.
  3. Enter your retail license key. You may need access to the Internet or to call Microsoft to complete the activation.
- **MAK volume licensing**—Multiple Activation Key (MAK) licensing allows the activation of multiple operating system installations using the same activation key.
  1. View or download the Microsoft Volume Activation Deployment Guide from the Microsoft web site.

2. Using an administrative user account, open a command prompt and follow the instructions from the deployment guide to activate MAK clients. Multiple reboots may be necessary before you can access a command prompt. You may need access to the Internet or to call Microsoft to complete the activation.

- **KMS volume licensing**—Key Management Service (KMS) licensing allows IT professionals to complete activations on their local network without contacting Microsoft.
    1. View or download the Microsoft Volume Activation Deployment Guide from the Microsoft web site.
    2. Using an administrative user account, open a command prompt and follow the instructions from the deployment guide to convert a MAK activation client to a KMS client. Multiple reboots may be necessary before you can access a command prompt.

Failing over from the command line

You can configure connections and initiate failover without using the Full-Server Failover Manager user interface. The same executable that launches the user interface can be used from a command prompt with options. The command line execution opens the user interface, passes through specified parameters, and initiates specified processes. You may want to use this alternate execution if you have different configuration files that you want to execute or if you have multiple connections. The Full-Server Failover Manager command can be initiated one at a time from a command prompt, or it can be scripted. The Full-Server Failover Manager executable is located in the installation directory.

Command

FFMANAGER

Description

Initiates the Full-Server Failover Manager user interface, passes through specified parameters, and initiates specified processes

Syntax

FFMANAGER /SOURCE ***source_name*** /TARGET ***target_name***/USERNAME ***username*** /PASSWORD ***password*** /VALIDATE /FIXALL /PROTECT /FAILOVER /LOGLEVEL ***number*** /CONFIG ***filename***

Options

- SOURCE ***source_name***—Name of the source
- TARGET ***target_name***—Name of the target
- USERNAME ***username***—Name of a user who is a member of the **Double-Take Admin** security group
- PASSWORD ***password***—Password associated with the specified user
- VALIDATE—Validates the configuration of the two servers to make sure they are compatible
- FIXALL—Corrects those errors that Full-Server Failover can automatically correct
- PROTECT—Initiates the connection between the source and target
- FAILOVER—Initiates failover from the source to the target

- LOGLEVEL number—Specifies the level of detailed logged based on the following numbers.
- 2—Informational messages are logged
- 3—Informational and error messages are logged
- 4—Informational, error, and exception messages are logged
- 5—Informational, error, exception, and debug messages are logged. This is the default setting.
- 6—Informational, error, exception, debug, and internal coding messages are logged
- CONFIG *filename*—Name of the file that contains the failover options. If no file is specified, the FFMDefaults.xml file will be used.

Examples

- ffmanager /source alpha /target beta /username administrator /password password /validate /fixall /protect
- ffmanager /source alpha /target beta /username administrator /password password /validate /failover

Notes

- If you do not specify any options with this command, the Full-Server Failover Manager user interface will open. The fields will be blank and no processing will occur.
- If you are failing over a cluster node, it is possible that volumes may lose their drive letter assignments. If a clustered application fails to start after failover and the disk signature has changed, check the drive letter assignments under the Disk Management utility and re-create drive letter assignments as needed.
- Because the Windows product activation is dependent on hardware, you may need to reactivate your Windows registration after failover. In most cases when you are using Windows 2003, you can follow the on-screen prompts to complete the reactivation. However, when you are using Windows 2008, the reactivation depends on your licensing type. If a Windows 2008 target comes online after failover with an activation failure, use the steps appropriate for your license type.
  - **Retail licensing**—Retail licensing allows the activation of a single operating system installation.
    1. Open the **System** applet in Windows **Control Panel**.
    2. Under **Windows activation** at the bottom of the page, click **Change product key**.

3. Enter your retail license key. You may need access to the Internet or to call Microsoft to complete the activation.

- **MAK volume licensing**—Multiple Activation Key (MAK) licensing allows the activation of multiple operating system installations using the same activation key.
  1. View or download the Microsoft Volume Activation Deployment Guide from the Microsoft web site.
  2. Using an administrative user account, open a command prompt and follow the instructions from the deployment guide to activate MAK clients. Multiple reboots may be necessary before you can access a command prompt. You may need access to the Internet or to call Microsoft to complete the activation.
- **KMS volume licensing**—Key Management Service (KMS) licensing allows IT professionals to complete activations on their local network without contacting Microsoft.
  1. View or download the Microsoft Volume Activation Deployment Guide from the Microsoft web site.
  2. Using an administrative user account, open a command prompt and follow the instructions from the deployment guide to convert a MAK activation client to a KMS client. Multiple reboots may be necessary before you can access a command prompt.

Failing over application workloads configured for DNS failover

When a failover condition has been met, failover will be triggered automatically if you configured automatic failover when establishing protection. If you configured manual intervention before failover, you can use the Application Manager to initiate failover for application workloads configured for DNS failover.

Note: In a clustered environment where the source suddenly becomes unavailable (for example, it crashes) and the Application Manager is open, it may appear to be unresponsive for up to 30 minutes before the failover process continues. The Application Manager is waiting on a Microsoft cluster file to respond. To reduce the amount of time before failover can continue, close and re-open the Application Manager.

If you are protecting a file server, failover is only available if the source is offline, in order to prevent name conflicts on the network.

---

1. If you are protecting a standalone BlackBerry server with Exchange or SQL, shutdown the BlackBerry server.
2. In the Application Manager, make sure your source target pair is selected and then on the **Monitor** tab, click **Failover**.

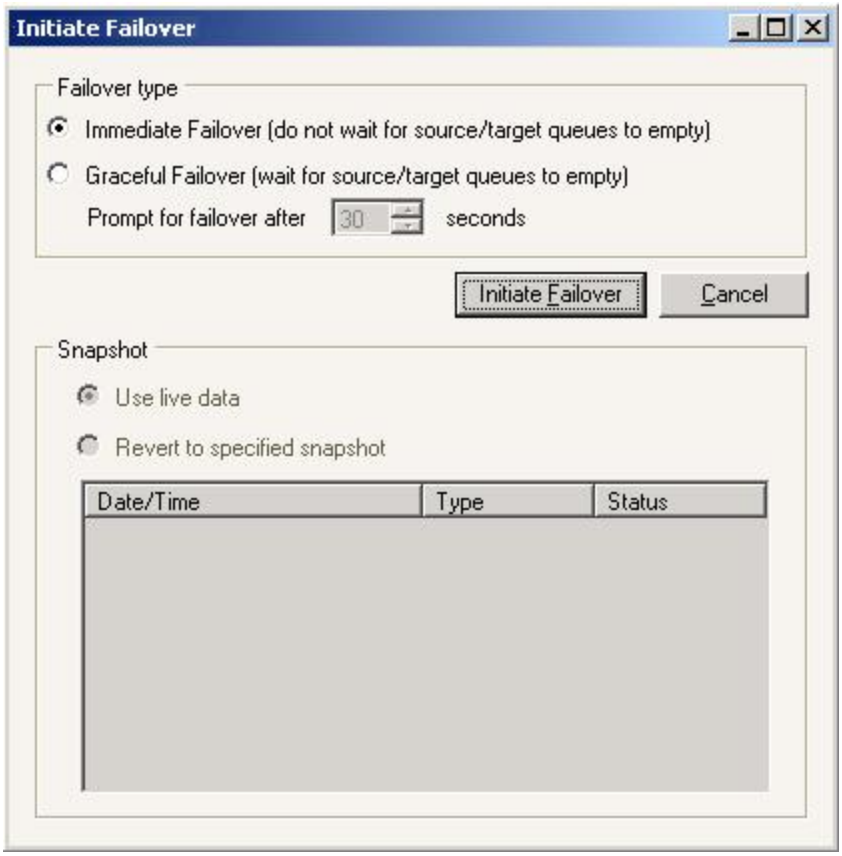

3. Specify if you want to perform an immediate or graceful failover. An immediate failover begins immediately without waiting for the data queues on the source and target to empty. A graceful failover waits until the queues are emptied before continuing. If you select the graceful failover, specify how often you want the failover prompt to continue asking for failover while there is still data in the queues.

Exchange Note: If you are protecting an Exchange virtual server in a cluster environment, you should use the graceful failover option so that the source cluster resources are taken offline gracefully.

SharePoint Note: The graceful failover option is not available.

File Server Note: The graceful failover option is not available.

---

4. If you have taken snapshots of your target data, specify the data you want to use for failover.
    - **Use live data**—Select this option to use the data on the target at the time of failover.
    - **Revert to specified snapshot**—Select this option, and then select a snapshot. The data on the target will be reverted to the selected snapshot. To help you understand what snapshots are available, use the **Type** and **Status** columns. The **Status** indicates the state of the connection between the source and target at the time the snapshot was taken. The **Type** indicates the kind of snapshot.
        - **Scheduled**—This snapshot was taken as part of a periodic snapshot.
        - **Deferred**—This snapshot was taken as part of a periodic snapshot, although it did not occur at the specified interval because the connection between the source and target was not in a good state.
        - **User Request**—This snapshot was taken manually by a user.
5. Click **Initiate Failover** to being the failover process.
6. If you shutdown your BlackBerry server, bring it back up when failover is complete and perform the following steps.
    a. Open a command prompt on the BlackBerry server to the \Program Files\Research In Motion\BlackBerry Enterprise Server\Utility directory.
    b. Run the following command.

        handheldcleanup –m

    c. Run the command, specifying the name of the BlackBerry server when prompted.

        handheldcleanup –u

    d. If you are using BlackBerry Enterprise Server 3.5 through 4.1, see the BlackBerry article How to move the BlackBerry Enterprise Server service account mailbox.

Exchange Note: Users using Outlook or Outlook Web Access to receive e-mail can connect after the changes have propagated through your environment. Users that had Outlook open during failover will need to restart the Outlook client (excluding Outlook Web Access clients on a LAN). Additionally, those users using Outlook Web Access or Outlook 2007 may see a security alert because the security certificate has the source server name but Exchange is now on the target. Click **Allow** or **OK** to dismiss the alert.

You will not be able to log in to the domain from the source Exchange server after failover because the target has assumed the source server's host Service Principal Name so that Outlook Web Access can use the source name. If you need to log in to the domain and Outlook Web Access is not needed, contact technical support for a workaround.

If you SMTP gateway is configured to send e-mail to a specific IP address that address is not failed over to the target, you will need to update the IP address after failover.

Mail stores or storage groups created after a failover will not be failed back.

SQL Note: After failover, linked databases in the SQL instance will be unavailable until the service master key is updated. You will need to run the command "alter service master key force regenerate" against the SQL target server to reset the service master key and then remove and re-add the linked servers into the target SQL instance.

After failover with a snapshot of a SQL database-only server, the SQL services on the target server are stopped and the databases are not mounted. You will need to manually start the MSSQLServer service for each instance on the target server and then manually attach the databases.

After failing over SQL 2008, Rich Internet Applications created using ADO.net 2 may not connect.

After failing over SQL 2008, you may not be able to take the SQL database offline. If this occurs, stop and restart the SQL Server Management Studio application, and then you can take the database offline.

After failing over in a SQL workgroup, you will not be able to connect to the source server instance of SQL. You can work around this issue by creating an alias on the target with the source server’s name.

BlackBerry Note: There is a potential to lose BlackBerry instant messages and instant message contacts when an Exchange mailbox is moved, depending on the type of BlackBerry hand-held device. You should back up

instant messages through the BlackBerry desktop software or have the BlackBerry Enterprise server configured to audit all instant messages through a policy.

Virtual workload failover

When a failover condition has been met, failover will be triggered automatically if you configured automatic failover when establishing protection. If you configured manual intervention before failover, you can failover your protection from the console you used to establish protection. If you are protecting host-level virtual disk files (the .vmdk files) from an ESX source to an ESX target, you will need to use the Double-Take Availability for VMware Infrastructure console to initiate failover. For all other virtual workloads, use the Double-Take Console to initiate failover.

Failing over virtual workloads in the Double-Take Console

1. On the **Monitor Connections** page, select the connection that you want to failover.
2. In the lower pane, click **Failover** in the toolbar.
3. Select the type of failover to perform.
   - **Live failover**—Select this option to initiate a full, live failover. This option will shutdown the source virtual machine (if available), stop the protection job, and start the replica virtual machine on the target with full network connectivity.
   - **Test failover**—Select this option to perform a test failover without network connectivity. This option will leave the source virtual machine online, suspend the protection job, and start the replica virtual machine on the target without network connectivity.
4. Click **Failover** to begin failover.

Note: Once failover has occurred, if you add CPUs to the replica of the source on the target, you may have to reboot the replica before the operating system will recognize the additional CPUs.

---

5. After failover is complete, you can undo it by clicking **Undo Failover** in the toolbar. The replica virtual machine on the target will be shutdown, the source virtual machine will be restarted (in the case of live failover), and protection will be restarted performing a file differences remirror. All changes made on the replica virtual machine on the target will be lost. If you do not want to lose changes made replica virtual machine on the target, perform a restore and failback.

Failing over virtual workloads in the Double-Take Availability for VMware Infrastructure console

1. Make sure your protection job is in an active (non-stopped) state.
2. On the **Monitor protection** page, select the connection that you want to failover and click **Failover** in the toolbar.
3. Select the type of failover to perform.
    - **Live failover**—Select this option to initiate a full, live failover. This option will shutdown the source virtual machine (if available), stop the protection job, and start the replica virtual machine on the target with full network connectivity.
    - **Test failover**—Select this option to perform a test failover without network connectivity. This option will leave the source virtual machine online, suspend the protection job, and start the replica virtual machine on the target without network connectivity.
4. Select the failover timing.
    - **Complete the current replication cycle before failover**—With this option, failover will begin immediately after the current replication cycle is completed. This option is not available if the job is in a stopped or error state.
    - **Failover immediately**—With this option, failover will begin immediately without waiting for the current replication cycle to complete.
5. Click **Failover** to begin failover.
6. After failover is complete, you can undo it by clicking **Undo Failover** in the toolbar. The replica virtual machine on the target will be shutdown, the source virtual machine will be restarted (in the case of live failover), and protection will be restarted performing a file differences remirror. All changes made on the replica virtual machine on the target will be lost.

Failback and restore

Your failover and restoration process will depend on the type of workload you are protecting.

- Failback for data workloads
- Failing back a full-server workload
- Failing back an application workload
- Failing back a virtual workload

Data workload failback and restoration

Failover occurred because the target was monitoring the source for a failure, and when a failure occurred, the target stood in for the source. User and application requests that were directed to the failed source are routed to the target.

While the users are accessing their data on the target, you can repair the issue(s) on the source. Before users can access the source again, you will need to restore the data from the target back to the source and perform failback. Failback is the process where the target releases the source identity it assumed during failover. Once failback is complete, user and application requests are no longer routed to the target, but back to the source.

Ideally, you want to restore your data from the target back to the source before you failback. This allows users who are currently accessing their data on the target because of failover to continue accessing their data. Restoration before failback reduces user downtime. The procedure to restore and then failback varies widely with server and network configuration. Another method, which may be easier in some environments, allows you to failback first and then restore the data from the target to the source. A possible disadvantage to this process is that users may experience longer downtime, depending on the amount of data to be restored, because they will be unable to access their data during both the restoration and the failback.

- Restoring then failing back
- Failing back then restoring

Restoring then failing back

Restoring before failing back allows your users to continue accessing their data on the failed over target, which is standing in for the source, while you perform the restoration process. The key to this process is to keep the users off of the source, but allow the source and target to communicate to perform the restoration.

1. Locate the file connect.sts on the source where you installed Double-Take Availability and rename it to connect.sts.old. This will keep the original connection from reconnecting when you bring the source online.
2. Resolve the problem(s) on the source that caused it to fail. Make sure in resolving the problems, that you do not bring the source on the network at this time because the target currently has the source's identity because of failover.
3. Disable all of the NICs on the source.
4. Change one of the NICs on the source to a unique IP address that the target can access.
5. Configure that IP address so that it does not automatically register with DNS. This option is on the Advanced TCP/IP Settings dialog box on the DNS tab.
6. Do not enable the modified NIC yet. If you do, you will receive a network name conflict, because the target has the source's identity because of failover. There are many variations for dealing with a name conflict, here are a few examples.
    - Enable the modified NIC, knowing you will get the name conflict error. Disregard the error. Change the source name to a unique name in the domain and reboot when prompted.
    - Change the source name to a unique name in a workgroup, not in the domain, and reboot when prompted. Enable the modified NIC.
    - Stop the Workstation and Server services on the source. You may be prompted to stop other services. Stop those services also and note the service names for later. Enable the modified NIC. The server will not broadcast its name to the network because of the services you disabled.
7. Stop any applications that may be running on your source. The files must be closed on the source so that updated files from the target will overwrite the files on the source.
8. At this point, confirm you have the following configuration.
    - Your target is standing in for your source because of failover, and users are accessing their data from the target.
    - Your source is back online with a unique IP address and no network name conflicts.
    - The source and target can communicate with each other.
    - All applications on the source are stopped.

9. From your target, confirm the Replication Console is communicating with the source using the new IP address.
    a. From the Replication Console on the target, right-click the source and select **Remove**.
    b. Depending on your configuration, the source may be automatically inserted back into the Replication Console. If it is not, select **Insert**, **Server**. Specify the source server and click **OK**.
10. Begin your restoration process.
    a. From the Replication Console, select **Tools**, **Restoration Manager**.

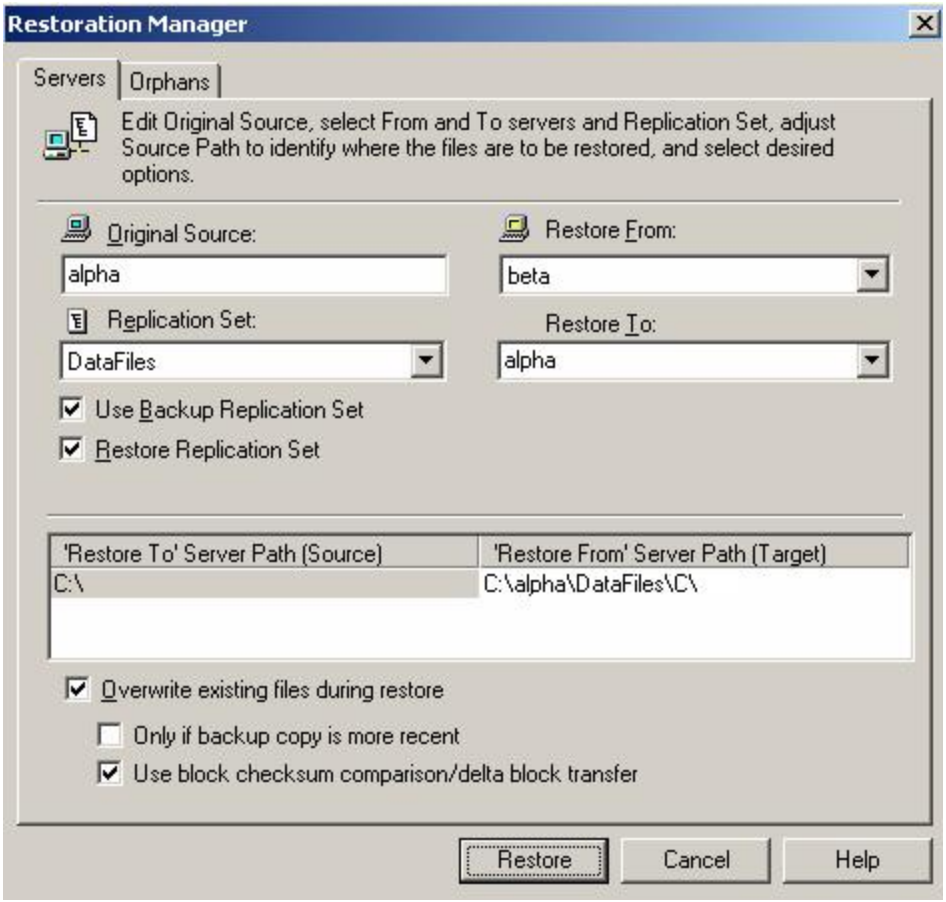

    b. Identify the **Original Source** machine. This is your source machine where the data originally resided.
    c. Select the **Restore From** machine. This is the target machine where the copy of the data is stored.

Note: If your target is a cluster, you can specify just one node in the cluster and restore only from that node. If you need to have the cluster

functionality (Microsoft failover from node to node) available during the restoration process, you will have to create a resource, like you did for your original connection, for the restoration process. See Protecting a standard cluster. Keep in mind when restoring, your target will function as your source, sending data to the target, which will be your original or new source.

---

d. **Replication Set** contains the replication set information stored on the target machine (the machine in **Restore From**). If no replication sets are available, the list will be blank. Select the replication set that corresponds to the data that you need to restore.
e. Select the **Restore To** machine. This is your temporary source that has the unique IP address.
f. The **Restore To** and **Restore From** paths will automatically be populated when the replication set is selected. The restore to path is the directory that is the common parent directory for all of the directories in the replication set. If the replication set crosses volumes, then there will be a separate path for each volume. The restore from path is the path on the target server where the replicated files are located.

Note: Restoring across a NAT router requires the ports to be the same as the original connection. If the ports have been modified (manually or reinstalled), you must set the port numbers to the same values as the last valid source/target connection.

---

g. Select the **Use Backup Replication Set** check box to use the target's copy of the replication set database for the restoration. If this check box is not marked, you will be accessing the replication set information from the source.
h. Select the **Restore Replication Set** check box to restore the target's copy of the replication set database to the source during the restoration process.
i. Select the restoration conditionals that you want to use.
   - **Overwrite existing files during restore**—This option restores all existing files by overwriting them. Any files that do not exist on the source are written also. If this option is disabled, only files that do not exist on the source will be restored.
   - **Only if backup copy is more recent**—This option restores only those files that are newer on the target than on the source. The entire file is

overwritten with this option.

Note: If you are using a database application, do not use the newer option unless you know for certain you need it. With database applications, it is critical that all files, not just some of them that might be newer, get mirrored.

- **Use block checksum comparison/delta block transfer**—Specify if you want the restoration process to use a block checksum comparison to determine which blocks are different. If this option is enabled, only those blocks (not the entire files) that are different will be restored to the source.

j. If you want to configure orphan files, click the **Orphans** tab. The same orphan options are available for a restoration connection as a standard connection.

k. If your original source was using **Replicate NT Security by Name**, you must enable that option on the target before you start the restoration. The option is available on the target's Server Properties on the Source tab.

l. Click **Restore** to begin the restoration. You can identify a restoration connection because it is enclosed in parenthesis ( ) and it has _Restore appended to the end of the replication set name. The initial restoration is complete when the **Mirror Status** is **Idle**. After the **Mirror Status** is **Idle**, the connection will continue replicating any on-going data changes from the target to the source.

11. After the **Mirror Status** is **Idle**, schedule a time for failback. User downtime will begin once failback is started, so select a time that will have minimal disruption on your users.
12. When you are ready, begin the failback process.
    a. Open the Failover Control Center.
    b. Select the target that is currently standing in for the failed source.
    c. Select the failed source and click **Failback**. The user downtime starts now. If you have a pre-failback script configured, it will be started.

Note: If the target is a cluster, you will need to determine the active node and failback from that node. Then you will need to failback from each of the other nodes to synchronize all of the nodes of the cluster.

---

d. When failback is complete, the post-failback script, if configured, will be started. When the script is complete, you will be prompted to determine if you want to continue monitoring the source, do not select either option. Leave the prompt dialog box open as is.

13. Back in the Replication Console, watch the restoration connection until activity has ended and replication is in a **Ready** state. This will happen as the final data in queue, if any, is applied on the source. The replication **Ready** state indicates replication is waiting for new incoming data changes. There will not be any additional data changes because failback has released the source identity, so users are no longer accessing their data.
14. Once replication is in a **Ready** state, disconnect the restoration connection from the target. This is the connection enclosed in parenthesis () and it has _Restore appended to the end of the replication set name.

Note: If your target is a cluster, take the Double-Take Source Connection resource offline to disconnect the connection.

During the restoration, only the data is restored back to the source. Shares are not created on the source during the restoration. Shares that were created on the target during failover will need to be created manually on the source.

---

15. On the source, change the IP address that you modified earlier to the unique address back to its original address. You can also enable any other NICs on the source.
16. Also on the source, change the source name back to its original name and reboot, or restart the Workstation, Server, and any other services you were prompted to stop.
17. Once the source is back online, users can reconnect to the source.
18. Confirm the Replication Console is communicating with the source using the original IP address.
    a. Right-click the source and select **Remove**.
    b. Depending on your configuration, the source may be automatically inserted back into the Replication Console. If it is not, select **Insert**, **Server**. Specify the source server and click **OK**.
19. At this time, you can go back to the dialog box in the Failover Control Center. Select **Continue** or **Stop** to indicate if you want to continue monitoring the source. After you have selected whether or not to continue monitoring the source, the

source post-failback script, if configured, will be started.

Note: The source must be online and Double-Take Availability must be running to ensure that the source post-failback script can be started. If the source has not completed its boot process, the command to start the script may be lost and the script will not be initiated.

---

20. You can now reconnect your original replication set on the source to your target, to reestablish protection.

Failing back then restoring

Failback before restoration can be a simpler process, but it may require additional downtime. The amount of downtime will depend on the amount of data to be restored. Users must be kept off of the source and target during this entire process.

1. Locate the file connect.sts on the source where you installed Double-Take Availability and rename it to connect.sts.old. This will keep the original connection from reconnecting when you bring the source online.
2. Resolve the problem(s) on the source that caused it to fail. Make sure in resolving the problems, that you do not bring the source on the network at this time because the target currently has the source's identity because of failover.
3. Schedule a time for this process. Select a time that will have minimal disruption on your users.
4. When you are ready, begin the failback process.
    a. Open the Failover Control Center and select the target that is currently standing in for the failed source.
    b. Select the failed source and click **Failback**. The user downtime starts now. If you have a pre-failback script configured, it will be started.

    Note: If the target is a cluster, you will need to determine the active node and failback from that node. Then you will need to failback from each of the other nodes to synchronize all of the nodes of the cluster.

    ---

    c. When failback is complete, the post-failback script, if configured, will be started. When the script is complete, you will be prompted to determine if you want to continue monitoring the source, do not select either option. Leave the prompt dialog box open as is.
5. Bring your source onto the network at this time, but make sure that the users are not accessing it. The target must be able to access the source, but users cannot access the source because the data on the source is out-of-date.
6. Once your source is on the network, select **Continue** or **Stop** in the Failover Control Center to indicate if you want to continue monitoring the source. After you have selected whether or not to continue monitoring the source machine, the source post-failback script, if configured, will be started.

Note: The source must be online and Double-Take Availability must be running to ensure that the source post-failback script can be started. If the source has not completed its boot process, the command to start the script may be lost and the script will not be initiated.

7. Stop any applications that may be running on your source. The files must be closed on the source so that updated files from the target will overwrite the files on the source.
8. Now you can begin your restoration process.
   a. From the Replication Console on the target, select **Tools**, **Restoration Manager**.

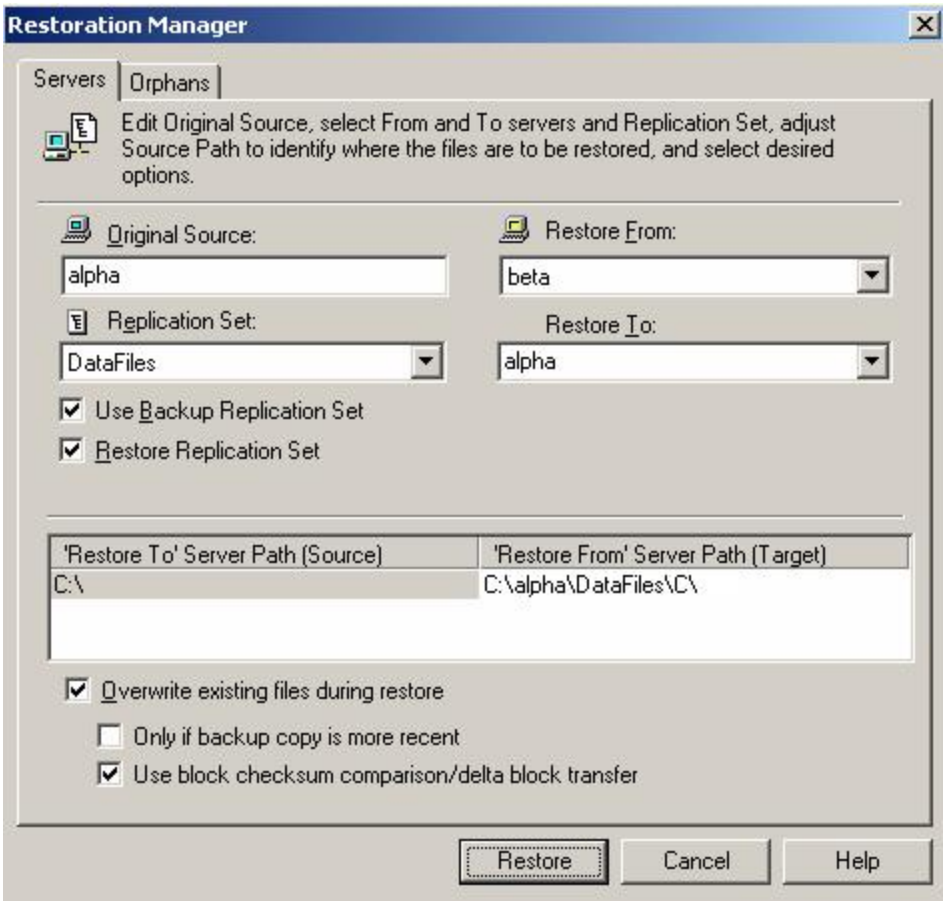

   b. Identify the **Original Source** machine. This is your source machine where the data originally resided.
   c. Select the **Restore From** machine. This is the target machine where the copy of the data is stored.

Note: If your target is a cluster, you can specify just one node in the cluster and restore only from that node. If you need to have the cluster functionality (Microsoft failover from node to node) available during the restoration process, you will have to create a resource, like you did for your original connection, for the restoration process. See Protecting a standard cluster. Keep in mind when restoring, your target will function as your source, sending data to the target, which will be your original or new source.

d. **Replication Set** contains the replication set information stored on the target machine (the machine in **Restore From**). If no replication sets are available, the list will be blank. Select the replication set that corresponds to the data that you need to restore.
e. Select the **Restore To** machine. This is your source where the updated data from the target will be sent.
f. The **Restore To** and **Restore From** paths will automatically be populated when the replication set is selected. The restore to path is the directory that is the common parent directory for all of the directories in the replication set. If the replication set crosses volumes, then there will be a separate path for each volume. The restore from path is the path on the target server where the replicated files are located.

Note: Restoring across a NAT router requires the ports to be the same as the original connection. If the ports have been modified (manually or reinstalled), you must set the port numbers to the same values as the last valid source/target connection.

g. Select the **Use Backup Replication Set** check box to use the target's copy of the replication set database for the restoration. If this check box is not marked, you will be accessing the replication set information from the source.
h. Select the **Restore Replication Set** check box to restore the target's copy of the replication set database to the source during the restoration process.
i. Select the restoration conditionals that you want to use.
   - **Overwrite existing files during restore**—This option restores all existing files by overwriting them. Any files that do not exist on the source are written also. If this option is disabled, only files that do not

exist on the source will be restored.

- **Only if backup copy is more recent**—This option restores only those files that are newer on the target than on the source. The entire file is overwritten with this option.

Note: If you are using a database application, do not use the newer option unless you know for certain you need it. With database applications, it is critical that all files, not just some of them that might be newer, get mirrored.

- **Use block checksum comparison/delta block transfer**—Specify if you want the restoration process to use a block checksum comparison to determine which blocks are different. If this option is enabled, only those blocks (not the entire files) that are different will be restored to the source.

j. If you want to configure orphan files, click the **Orphans** tab. The same orphan options are available for a restoration connection as a standard connection.

k. If your original source was using **Replicate NT Security by Name**, you must enable that option on the target before you start the restoration. The option is available on the target's **Server Properties** on the **Source** tab.

l. Click **Restore** to begin the restoration. You can identify a restoration connection because it is enclosed in parenthesis ( ) and it has _Restore appended to the end of the replication set name.

Note: If your target is a cluster, take the Double-Take Source Connection resource offline to disconnect the connection.

During the restoration, only the data is restored back to the source. Shares are not created on the source during the restoration. Shares that were created on the target during failover will need to be created manually on the source.

---

9. Because there are no users accessing the target data, the restoration process is complete when the **Mirror Status** is **Idle**. When the **Mirror Status** is **Idle**, disconnect the restoration connection from the target.

10. Your original connection on the source, if it still exists in the Replication Console, will be in a **Mirror Required** state. Right-click the connection and select **Mirror**, **Start**. Select the type of mirror you wish to perform and click **OK**. When prompted to start replication, click **Yes**.
11. Once you have restarted the mirror and replication, you can allow users to reconnect to the source.

Failing back and restoring a full-server workload

After your target has failed over and becomes your source, you can stay with that configuration long term. However, in some instances, it may be necessary or desired to go back to using the original hardware after you have failed over. Use the following process to failback to your original (or other) hardware.

1. Because your new source is on the network, you must make sure your original source is unique on the network to avoid name and IP address conflicts. You have several options available for achieving this. .
    - Reinstall Windows using unique server information. This may be the best option if your original source was a domain controller or running a name-specific application like Exchange.
    - Run a utility like Microsoft SysPrep to modify SIDs (security identifiers), IP addresses, and the server name. This option is only available for standalone servers. You cannot run SysPrep on a computer that has been configured as a Cluster Service server, a Certificate Services server, or a domain controller.
    - Manually make the original source unique by modifying IP addresses and the server name. If your original source was a domain controller, you must also modify the SIDs.
2. Establish protection from your new source (the server that you failed over to) back to your original source using the same process as when you protected your original source. The mirror that completes from the new source back to the original source will act as a restoration process.
3. When the protection is established and the mirror is complete, initiate failover. This process will act as a failback. Once the failover is complete, you will be back to your original source.

Application failback and restoration

When protecting application workloads, your failback and restoration process will depend on if you configured your application for identity or DNS failover.

- **Identity failover**—If your application is configured for identity failover, you will need to stop your application services and then you have the choice of restoring and then failing back or failing back and then restoring. See Failback and restoration for applications configured for identity failover.
- **DNS failover**—If your application is configured for DNS failover, you will perform your restoration and then failback. See Restoring then failing back applications configured for DNS failover.

Failback and restoration for applications configured for identity failover

For application workloads that were configured for identity failover, you need to stop the services on your source that correspond with your application before you begin the failback and restoration process. Use the table below as a guideline for the services to stop. After you have stopped the services on the source, use the instructions for data workload failback and restoration to complete your application failback and recovery.

Exchange Note: After the restoration is complete, you will need to rehome the informational store databases to the source.

1. From a command prompt on the source, run the post_restore_<source server name>_<target server name>.bat file that Application Manager automatically generated.
2. Restart any Outlook clients so that they can access the source.

Exchange 2007

- Microsoft Exchange Active Directory Topology Service
- Microsoft Exchange Anti-spam Update
- Microsoft Exchange EdgeSync
- Microsoft Exchange File Distribution
- Microsoft Exchange IMAP4
- Microsoft Exchange Information Store
- Microsoft Exchange Mail Submission
- Microsoft Exchange Mailbox Assistants
- Microsoft Exchange POP3
- Microsoft Exchange Replication Service
- Microsoft Exchange Search Indexer
- Microsoft Exchange Service Host
- Microsoft Exchange System Attendant
- Microsoft Exchange Transport
- Microsoft Exchange Transport Log Search
- Microsoft Search (Exchange)
- World Wide Web Publishing Service

Exchange 2003

- MSExchangeSA
- MSExchangeMGMT
- POP3SVC
- IMAP4SVC
- ResVC
- MSExchangeES
- W3SVC
- SMTPSVC

SQL Server

- MSSqlServer
- SQLServerAgent
- MSSearch (SQL 2000)
- MSFteSQL (SQL 2005)
- MSSQLServerADHelper
- MSDTC
- MSSQLServerOLAPService
- MSDTSServer
- SQLWriter
- SQLBrowser (SQL 2005)

SharePoint

- IISADMIN
- HTTPFilter
- SMTPSVC
- W3SVC
- SPAdmin
- SPSearch
- SPTimerV3
- SPTrace
- SPWriter

BlackBerry

- BlackBerry Router
- BlackBerry Server Alert
- BBAttachServer

- BlackBerry Controller
- BlackBerry Database Consistency Service
- BlackBerry Dispatcher
- BlackBerry Policy Service
- BlackBerry SyncServer
- BlackBerry MDS Connection Service
- MdsTomcat

File Server

- Server
- Computer Browser

Restoring then failing back applications configured for DNS failover

For application workloads that were configured for DNS failover, you can use the Application Manager to initiate failback and, if desired restore data from the target back to the source. In order to minimize downtime, the restoration process is completed before the failback.

1. If you are using a cluster environment and protecting Exchange, make sure the Physical Disk resource(s) and the IP Address resource are online on the source cluster.

   Exchange Note: When you bring the source cluster online, an identical network name will be active on the target. Therefore, when the source cluster tries to bring the Exchange virtual server on the source online, the network name resource will fail and the group will not come online. Allow the source cluster to finish trying to bring the resources online before beginning failback.

2. When you bring the source cluster online, an identical network name will still be active on the target. Because of this, when the source cluster tries to bring up the EVS on the source, the network name resource will fail and consequently the group will not come online on the source. You should allow the source cluster to finish trying to bring the resources online before using the to failback.
3. In the Application Manager, make sure your source target pair is selected and then on the **Monitor** tab, click **Failback**.

   Note: The **Failback** button may not become active right away after completing a failover. In this case, restart the Application Manager.

4. Specify the options for your failback and restoration.

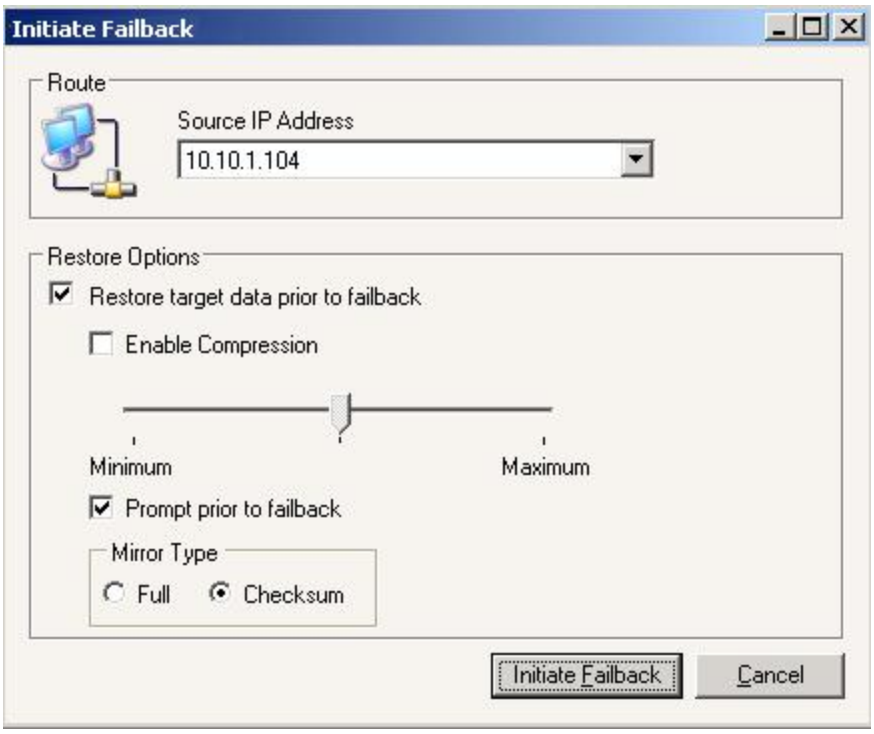

- **Source IP Address**—Select an IP address on the source to handle the restoration data. In a cluster environment running Exchange, select the name of Exchange virtual server dependent IP address.
- **Restore target data prior to failback**—Select this option to restore data from the target back to the source. If you are certain that there is no data that you want to restore or you are willing to lose any data changes on the target, you can disable this option.
- **Enable Compression**—Compression allows you to reduce the amount of bandwidth needed to transmit restoration data from the target back to the source. The data is compressed before being transmitted and then is uncompressed before it is written on the source. Typically, compression is used in WAN environments, but not in LAN environments. If desired, enable compression and select the level of compression that you want to use.
- **Prompt prior to failback**—When this option is enabled, the failback process will not start until you manually initiate it after the restoration is complete. You will be prompted when the restoration process is complete. If you disable the failback prompt, failback will automatically start when the restoration is complete.
- **Mirror**—Select the type of mirror to use for the restoration process. A **Full** mirror will transmit all files from the target back to the target. A **Checksum** mirror will transmit only the blocks of data that are different between the target and source.

5. Click **Initiate Failback** to being the failback process.

Exchange Note: If you deslected any mail stores during your failover configuration, you may see a message during failback about potential errors (unpaired mail stores). This message can be disregarded.

If you created any new mail stores on the target after failover, they will not be failed back.

Mail sent to public folders during failback may be routed to the target server after Exchange is shut down, which will result in mail being stuck in the queue. Make sure mail is not sent to public folders until the failback process is complete.

If you close the Application Manager during failback, you may have to manually run the post_restore.bat file which starts the Exchange services and updates public folders on the source.

If your source is an Exchange 2007 CCR, LCR, or SCC cluster, after failback the CCR, LCR, or SCC replication will need to be manually reseeded after verifying Exchange is functioning properly. For information about this process, see Microsoft TechNet Article How to Seed a Cluster Continuous Replication Copy.

In a like-named cluster environment with more than one DNS server, there may be a delay contacting the source after failback. The DNS server used by the source cluster is updated on failback to point back to the source server. However, if the Application Manager is running on a machine that uses a different DNS server, it may not recognize the change until the next DNS zone refresh.

Restoring then failing back virtual workloads

1. Open the console you used to establish protection.
2. On the **Monitor Connections** or **Monitor protections** page (depending on your console), select the connection that you want to restore and failback.
3. Click **Reverse Protection** in the toolbar. The flow of mirroring and replication data will change. Data will be transmitted from the replica virtual machine on the target back to the source.
4. After mirroring is complete and you are ready to failback, perform the failover process again. This process will act as a failback. Once the failover is complete, you will be back to your original source.

Connections

A unique connection ID is associated with each Double-Take Availability connection. The connection ID provides a reference point for each connection. The connection ID is determined by sequential numbers starting at one (1). Each time a connection is established, the ID counter is incremented. It is reset back to one each time the Double-Take service is restarted. For example, if the Double-Take service was started and the same replication set was connected to five target machines, each connection would have a unique connection ID from 1 to 5. The connection can be in various states.

- **Started**—The network connection exists and is available for data transmission. Replication and mirror data are transmitted to the target as soon as possible. This is the standard state that you will see most often.
- **Stopped**—Double-Take Availability has linked the source and target, but the network connection does not exist. Replication and mirror data are not transmitted to the target but are held in queue on the source.
- **Paused**—The network connection exists and is available for data transmission, but the replication and mirror data is being held in a queue and is not being transmitted to the target.
- **Scheduled**—Double-Take Availability has linked the source and target, but the network connection is not established until event driven or scheduling criteria have been met.
- **Error**— A transmission error has occurred. Possible errors include a broken physical line or a failed target service.

You can perform the following functions to manage your connections.

- Data queues
- Auto-disconnect and auto-reconnect
- Pausing and resuming target processing
- Blocking writing to the target paths
- Disconnecting a connection

Data queues

During the Double-Take Availability installation, you identified the amount of disk space that can be used for Double-Take Availability queuing. Queuing to disk allows Double-Take Availability to accommodate high volume processing that might otherwise fill-up system memory. For example, on the source, this may occur if the data is changing faster than it can be transmitted to the target, or on the target, a locked file might cause processing to back up.

The following diagram will help you understand how queuing works. Each numbered step is described after the diagram.

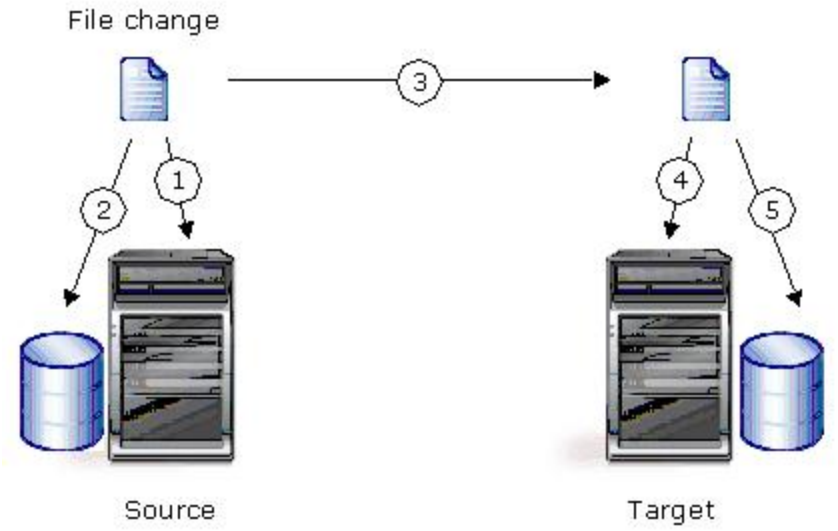

1. If data cannot immediately be transmitted to the target, it is stored, or queued, in system memory. You can configure how much system memory you want to use for queuing. By default, 128 or 512 MB of memory is used, depending on your operating system.
2. When the allocated amount of system memory is full, new changed data bypasses the full system memory and is queued directly to disk. Data queued to disk is written to a transaction log. Each transaction log can store 5 MB worth of data. Once the log file limit has been reached, a new transaction log is created. The logs can be distinguished by the file name which includes the target IP address, the Double-Take Availability port, the connection ID, and an incrementing sequence number.

Note: You may notice transaction log files that are not the defined size limit. This is because data operations are not split. For example, if a transaction log

has 10 KB left until the limit and the next operation to be applied to that file is greater than 10 KB, a new transaction log file will be created to store that next operation. Also, if one operation is larger than the defined size limit, the entire operation will be written to one transaction log.

---

3. When system memory is full, the most recent changed data is added to the disk queue, as described in step 2. This means that system memory contains the oldest data. Therefore, when data is transmitted to the target, Double-Take Availability pulls the data from system memory and sends it. This ensures that the data is transmitted to the target in the same order it was changed on the source. Double-Take Availability automatically reads operations from the oldest transaction log file into system memory. As a transaction log is depleted, it is deleted. When all of the transaction log files are deleted, data is again written directly to system memory (step 1).
4. To ensure the integrity of the data on the target, the information must be applied in the same order as it was on the source. If there are any delays in processing, for example because of a locked file, a similar queuing process occurs on the target. Data that cannot immediately be applied is queued to system memory. By default, 128 or 512 MB of memory is used, depending on your operating system.
5. When the allocated amount of system memory on the target is full, new incoming data bypasses the full system memory and is queued directly to disk. Data queued to disk is written to a transaction log. On the target, the transaction logs are identified with the source IP address, the Double-Take Availability port, the connection ID, and an incrementing sequence number.

Like the source, system memory on the target contains the oldest data so when data is applied to the target, Double-Take Availability pulls the data from system memory. Double-Take Availability automatically moves operations from the oldest transaction log file to system memory. As a transaction log is depleted, it is deleted. When all of the transaction log files are deleted, data is again written directly to system memory (step 4).

Queuing data

You should configure queuing on both the source and target.

1. Open the Replication Console and right-click the server on the left pane of the Replication Console.
2. Select **Properties**.
3. Select the **Queue** tab.
4. Specify the queue settings for the server.

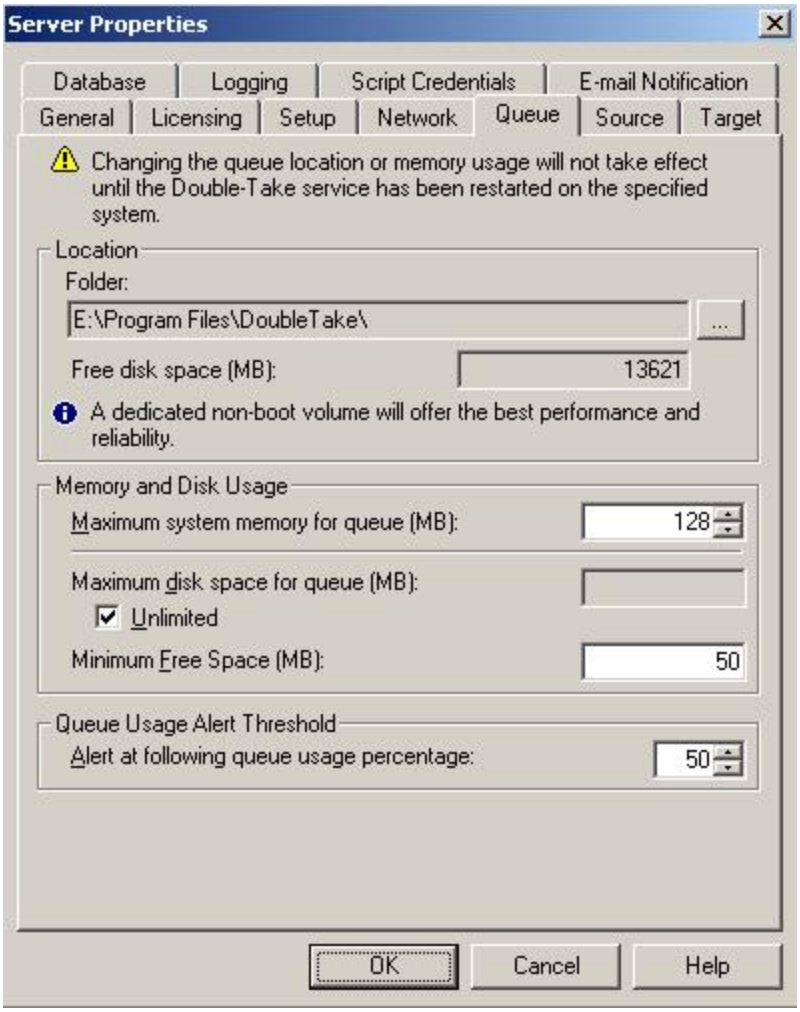

- **Folder**—This is the location where the disk queue will be stored. Double-Take Availability displays the amount of free space on the volume selected. Any changes made to the queue location will not take effect until the Double-Take service has been restarted on the server.

  When selecting the queue location, keep in mind the following caveats.

- Select a location on a non-clustered volume that will have minimal impact on the operating system and applications being protected.
- Select a location that is on a different volume as the location of the Windows pagefile.
- Select a dedicated, non-boot volume.
- Do not select the root of a volume.
- Do not select the same physical or logical volume as the data being replicated.

Although the read/write ratio on queue files will be 1:1, optimizing the disk for write activity will benefit performance because the writes will typically be occurring when the server is under a high load, and more reads will be occurring after the load is reduced. Accordingly, use a standalone disk, mirrored (RAID 1) or non-parity striped (RAID 0) RAID set, and allocate more I/O adapter cache memory to writes for best performance. A RAID 5 array will not perform as well as a mirrored or non-parity striped set because writing to a RAID 5 array incurs the overhead of generating and writing parity data. RAID 5 write performance can be up to 50% less than the write performance of a single disk, depending on the adapter and disk.

Another option is to use a solid state disk, which are hard drives that use RAM instead of disk platters. These devices are typically quite costly, but they will provide superior performance as a queuing device when the best performance is required.

Note: Scanning the Double-Take Availability queue files for viruses can cause unexpected results. If anti-virus software detects a virus in a queue file and deletes or moves it, data integrity on the target cannot be guaranteed. As long as you have your anti-virus software configured to protect the actual production data, the anti-virus software can clean, delete, or move an infected file and the clean, delete, or move will be replicated to the target. This will keep the target from becoming infected and will not impact the Double-Take Availability queues.

- **Maximum system memory for queue**—This is the amount of Windows system memory, in MB, that will be used to store data in queues. When exceeded, queuing to disk will be triggered. This value is dependent on the amount of physical memory available but has a minimum of 32 MB. By default, 128 MB or 512 MB of memory is used, depending on your operating

system. If you set it lower, Double-Take Availability will use less system memory, but you will queue to disk sooner which may impact system performance. If you set it higher, Double-Take Availability will maximize system performance by not queuing to disk as soon, but the system may have to swap the memory to disk if the system memory is not available.

Since the source is typically running a production application, it is important that the amount of memory Double-Take Availability and the other applications use does not exceed the amount of RAM in the system. If the applications are configured to use more memory than there is RAM, the system will begin to swap pages of memory to disk and the system performance will degrade. For example, by default an application may be configured to use all of the available system memory when needed, and this may happen during high-load operations. These high-load operations cause Double-Take Availability to need memory to queue the data being changed by the application. In this case, you would need to configure the applications so that they collectively do not exceed the amount of RAM on the server. Perhaps on a server with 1 GB of RAM running the application and Double-Take Availability, you might configure the application to use 512 MB and Double-Take Availability to use 256 MB, leaving 256 MB for the operating system and other applications on the system. Many server applications default to using all available system memory, so it is important to check and configure applications appropriately, particularly on high-capacity servers.

Any changes to the memory usage will not take effect until the Double-Take service has been restarted on the server.

- **Maximum disk space for queue**—This is the maximum amount of disk space, in MB, in the specified **Folder** that can be used for Double-Take Availability disk queuing, or you can select **Unlimited** which will allow the queue usage to automatically expand whenever the available disk space expands. When the disk space limit is reached, Double-Take Availability will automatically begin the auto-disconnect process. By default, Double-Take Availability will use an unlimited amount of disk space. Setting this value to zero (0) disables disk queuing.
- **Minimum Free Space**—This is the minimum amount of disk space in the specified **Folder** that must be available at all times. By default, 50 MB of disk space will always remain free. The **Minimum Free Space** should be less than the amount of physical disk space minus **Maximum disk space for queue**.

Note: The **Maximum disk space for queue** and **Minimum Free Space** settings work in conjunction with each other. For example, assume your queues are stored on a 10 GB disk with the **Maximum disk space** for queue set to 10 GB and the **Minimum Free Space** set to 500 MB. If another program uses 5 GB, Double-Take Availability will only be able to use 4.5 GB so that 500 MB remains free.

- **Alert at following queue usage percentage**—This is the percentage of the disk queue that must be in use to trigger an alert message in the Linux system log. By default, the alert will be generated when the queue reaches 50%.

5. Click **OK** to save the settings.

Auto-disconnect and auto-reconnect

While disk queues are user configurable and can be extensive, they are limited. If the amount of disk space specified for disk queuing is met, additional data could not be added to the queue and data would be lost. To avoid any data loss, the auto-disconnect and auto-reconnect processes occur.

- **Exhausted queues on the source**—If disk queuing is exhausted on the source, Double-Take Availability will automatically start disconnecting connections. This is called auto-disconnect. The transaction logs and system memory are flushed allowing Double-Take Availability to begin processing anew. The auto-reconnect process ensures that any connections that were auto-disconnected are automatically reconnected. Then, if configured, Double-Take Availability will automatically remirror the data. This process is called auto-remirror. The remirror re-establishes the target baseline to ensure data integrity, so disabling auto-remirror is not advised.
- **Exhausted queues on the target**—If disk queuing is exhausted on the target, the target instructs the source to pause. The source will automatically stop transmitting data to the target and will queue the data changes. When the target recovers, it will automatically tell the source to resume sending data. If the target does not recover by the time the source queues are exhausted, the source will auto-disconnect as described above. The transaction logs and system memory from the source will be flushed then Double-Take Availability will auto-reconnect. If configured, Double-Take Availability will auto-remirror. The remirror re-establishes the target baseline to ensure data integrity, so disabling auto-remirror is not advised.
- **Queuing errors**—If there are errors during disk queuing on either the source or target, for example, Double-Take Availability cannot read from or write to the transaction log file, the data integrity cannot be guaranteed. To prevent any loss of data, the source will auto-disconnect and auto-reconnect. If configured, Double-Take Availability will auto-remirror. The remirror re-establishes the target baseline to ensure data integrity, so disabling auto-remirror is not advised.
- **Target server interruption**—If a target machine experiences an interruption (such as a cable or NIC failure), the source/target network connection is physically broken but both the source and target maintain the connection information. The Double-Take Availability source, not being able to communicate with the Double-Take Availability target, stops transmitting data to the target and queues the data changes, similar to the exhausted target queues described above. When the interruption is resolved and the physical source/target connection is reestablished, the source begins sending the queued data to the target. If the source/target connection is not reestablished by the time the source queues are exhausted, the source will auto-disconnect as described above.

- **Target service shutdown**—If the target service is stopped and restarted, there could have been data in the target queue when the service was stopped. To prevent any loss of data, the Double-Take service will attempt to persist to disk important target connection information (such as the source and target IP addresses for the connection, various target queue information, the last acknowledged operation, data in memory moved to disk, and so on) before the service is stopped. If Double-Take Availability is able to successfully persist this information, when the Double-Take service on the target is restarted, Double-Take Availability will pick up where it left off, without requiring an auto-disconnect, auto-reconnect, or auto-remirror. If Double-Take Availability cannot successfully persist this information prior to the restart (for example, a server crash or power failure where the target service cannot shutdown gracefully), the source will auto-reconnect when the target is available, and if configured, Double-Take Availability will auto-remirror. The remirror re-establishes the target baseline to ensure data integrity, so disabling auto-remirror is not advised.

Note: If you are experiencing frequent auto-disconnects, you may want to increase the amount of disk space on the volume where the Double-Take Availability queue is located or move the disk queue to a larger volume.

If you have changed data on the target while not failed over, for example if you were testing data on the target, Double-Take Availability is unaware of the target data changes. You must manually remirror your data from the source to the target, overwriting the target data changes that you caused, to ensure data integrity between your source and target.

Reconnecting automatically

Use the following steps to configure automatic reconnections.

1. Open the Replication Console and right-click the server on the left pane of the Replication Console.
2. Select **Properties** and select the **Setup** tab.

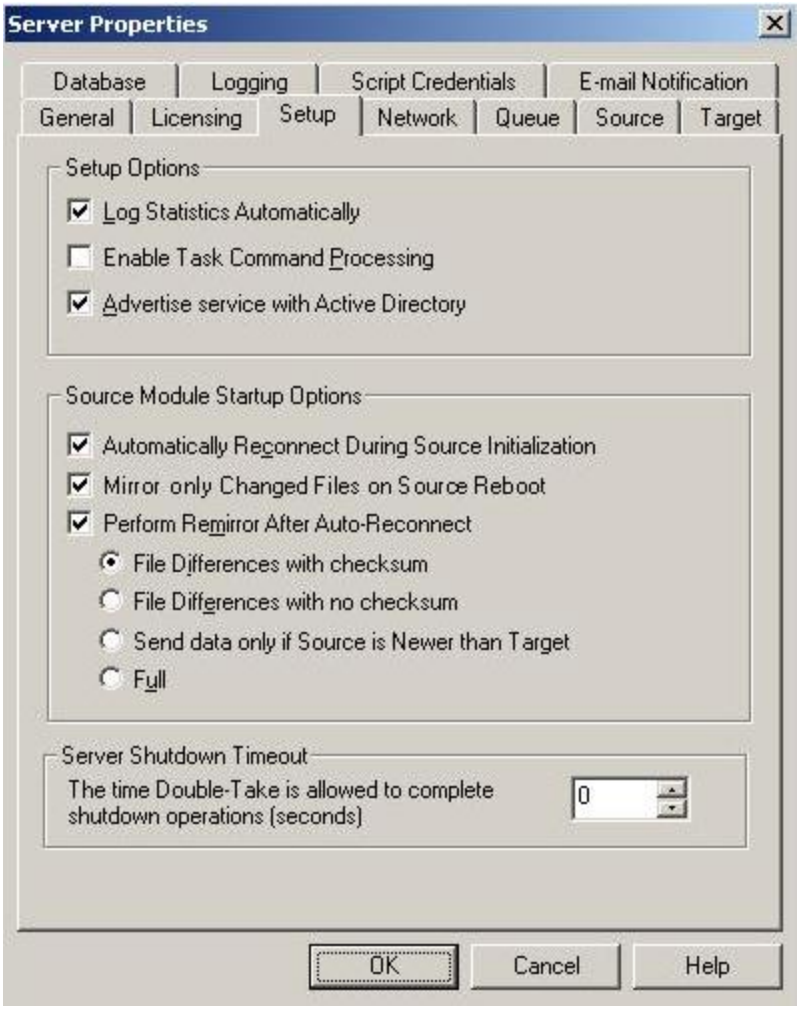

3. Verify that the check box **Automatically Reconnect During Source Initialization** is marked to enable the auto-reconnect feature.
4. Click **OK** to save the settings.

Pausing and resuming target processing

You can break the source/target connection without disconnecting the connection, so that you can control the transmission of data across the network. You can do this by pausing the target. If the target is paused, data is queued on the source until you manually resume the target.

For example, you must pause the target while you perform a backup of database files stored on the target because the database and log files must be backed up when they are at the exact same point in time. For example, say the backup of the file mydatabase.mdf begins on the target. While the backup program has access to the file, Double-Take Availability cannot write to the file. When the backup completes, Double-Take Availability writes to the file. Double-Take Availability also writes to the corresponding mydatabase.ldf file. When the backup gets to the mydatabase.ldf file, it no longer matches the .mdf file. The database may require special startup procedures, which may result in incomplete transactions being left in the database or data loss. To workaround this scenario, pause the target before starting the backup and then resume the target when the backup is complete.

While the target is paused, the Double-Take Availability source cannot queue data indefinitely. If the source queue is filled, Double-Take Availability will automatically disconnect the connections and attempt to reconnect them.

To pause a target, open the Replication Console and right-click a target server on the left pane of the Replication Console. Select **Pause Target**. All active connections to that target will complete the operations already in progress. You will see **Pause Pending** in the Replication Console while these operations are completed. The status will update to **Paused** after the operations are completed. Any new operations will be queued on the source until the target is resumed. When you are ready to resume the target, right-click the target and select **Resume Target**.

Note: If you have multiple connections to the same target, all connections will be paused and resumed.

Blocking writing to the target paths

You can block writing to the paths on the target that contain the copy of the replication set data. This keeps the data from being changed outside of Double-Take Availability processing. To block the replication set data paths on the target, open the Replication Console, and right-click the connection on the right pane of the Replication Console. Select **Block Target Path(s)**. To unblock the paths, right-click the connection and deselect **Block Target Path(s)**.

Note: If you are going to use failover, any target paths that are blocked will automatically be unblocked during the failover process so that users can modify data on the target after failover. During a restoration, the paths are automatically blocked again. If you failover and failback without performing a restoration, the target paths will remain unblocked. You can manually block or unblock the target paths by right-clicking on a connection.

Do not block your target paths if you are protecting a full-server because system state data will not be able to be written to the target.

Disconnecting a connection

To disconnect a Double-Take Availability connection, open the Replication Console, and right-click the connection on the right pane of the Replication Console. Select **Disconnect**. The source and target will be disconnected.

Note: If a connection is disconnected and the target is monitoring the source for failover, you will be prompted if you would like to continue monitoring for a failure. If you select **Yes**, the Double-Take Availability connection will be disconnected, but the target will continue monitoring the source. To make modifications to the failure monitoring, you will need to use the Failover Control Center. If you select **No**, the Double-Take Availability connection will be disconnected, and the source will no longer be monitored for failure by the target.

If a connection is disconnected while large amounts of data still remain in queue, the Replication Console may become unresponsive while the data is being flushed. The Replication Console will respond when all of the data has been flushed from the queue.

Mirroring

Mirroring is one of the key components of Double-Take Availability. You can perform the following functions to manage mirroring.

- Stopping, starting, pausing, or resuming mirroring
- Mirroring automatically
- Running scripts during mirroring
- Removing orphan files

Stopping, starting, pausing, or resuming mirroring

After a connection is established, you need to be able to control the mirroring. You can start, stop, pause and resume mirroring. open the Replication Console, and right-click the connection on the right pane of the Replication Console. Select **Mirroring** and the appropriate mirror control.

- **Pause or Resume**—When pausing a mirror, Double-Take Availability stops queuing mirror data on the source but maintains a pointer to determine what information still needs to be mirrored to the target. Therefore, when resuming a paused mirror, the process continues where it left off.
- **Stop**—When stopping a mirror, Double-Take Availability stops queuing mirror data on the source and does not maintain a pointer to determine what information still needs to be mirrored to the target. Therefore, when starting a mirror that has been stopped, the process will mirror all of the data contained in the replication set.
- **Start**—If you select to start a mirror, you will need to make the following two selections on the Start Mirror dialog box.
  - **Full Mirror**—All files in the replication set will be sent from the source to the target.
  - **File differences**—Only those files that are different based size or date and time will be sent from the source to the target. Expand *File difference mirror options compared* below to see how the file difference mirror settings work together, as well as how they work with the global checksum setting on the **Source** tab of the Server Properties.
  - **Send data only if Source is newer than Target**—Only those files that are newer on the source are sent to the target.

    Note: If you are using a database application, do not use the newer option unless you know for certain you need it. With database applications, it is critical that all files, not just some of them that might be newer, get mirrored.

  - **Use block checksum**—For those files flagged as different, the mirror performs a checksum comparison and only sends those blocks that are different.
  - **Calculate Replication Set size prior to mirror**—Determines the size of the replication set prior to starting the mirror. The mirroring status will update the percentage complete if the replication set size is calculated.

File difference mirror options compared

- **File Differences**—Any file that is different on the source and target based on the date, time, and/or size is transmitted to the target. The mirror sends the entire file.
- **File Differences and Only if Source is Newer**—Any file that is newer on the source than on the target based on date and/or time is transmitted to the target. The mirror sends the entire file.
- **File Differences and Checksum**—This option is dependent on the global checksum all option on the Server Properties source tab.
    - **Checksum All disabled**— Any file that is different on the source and target based on date, time, and/or size is flagged as different. The mirror then performs a checksum comparison on the flagged files and only sends those blocks that are different.
    - **Checksum All enabled**—The mirror performs a checksum comparison on all files and only sends those blocks that are different.
- **File Differences, Only if Source is Newer, and Checksum**—Any file that is newer on the source than on the target based on date and/or time is flagged as different. The mirror then performs a checksum comparison on the flagged files and only sends those blocks that are different.

Mirroring automatically

In certain circumstances, for example if the disk-based queues on the source are exhausted, Double-Take Availability will automatically disconnect connections (called auto-disconnect) and then automatically reconnect them (called auto-reconnect). In order to ensure data integrity on the target, Double-Take Availability will perform an automatic mirror (called an auto-remirror) after an auto-reconnect.

Note: Auto-remirror is a per source option. When enabled, all connections from the source will perform an auto-remirror after an auto-reconnect. When disabled, none of the connections from the source will perform an auto-remirror after an auto-reconnect.

---

1. Open the Replication Console, and right-click a server in the left pane of the Replication Console and select **Properties**.
2. Select the **Setup** tab.
3. Verify that **Perform Remirror After Auto-Reconnect** is selected to initiate an auto-remirror after an auto-reconnect.

   Note: If auto-remirror is disabled and an auto-reconnect occurs, the transmission state of the connection will remain pending after the reconnect until a mirror is started manually.

   ---

4. Verify that **Mirror only Changed Files on Source Reboot** is selected to use the Windows NTFS change journal to track file changes. When this option is enabled, if the source is rebooted, only the files identified in the change journal will be remirrored to the target. This setting helps improve mirror times.
5. Specify the type of mirror that you wish to perform.
   - **Differences with Checksum**—Any file that is different on the source and target based on date, time, and/or size is flagged as different. The mirror then performs a checksum comparison on the flagged files and only sends those blocks that are different.
   - **Differences with no Checksum**—Any file that is different on the source and target based on date, time, and/or size is sent to the target.
   - **Send data only if Source is newer than Target**—Only those files that are newer on the source are sent to the target.

- **Full**—All files are sent to the target.

Note: Database applications may update files without changing the date, time, or file size. Therefore, if you are using database applications, you should use the Differences with checksum or Full option.

Stopping, starting, pausing, or resuming mirroring contains a comparison of how the file difference remirror settings work together, as well as how they work with the global checksum setting on the **Source** tab of the Server Properties.

---

6. Click **OK** to save the settings.

Running scripts during mirroring

You can customize your mirroring process by running customized scripts on the target at predefined points in the mirroring process. Scripts may contain any valid Windows command, executable, batch, or script file. The scripts are processed using the same account running the Double-Take service, unless you provide specific credentials on the Server Properties Script Credentials tab for the target server. There are three types of mirroring scripts.

- **Mirror start**—This script starts when the target receives the first mirror operation. In the case of a difference mirror, this may be a long time after the mirror is started because the script does not start until the first different data is received on the target. If the data is synchronized and a difference mirror finds nothing to mirror, the script will not be executed.
- **Mirror stop**—This script starts when a mirror is stopped, which may be caused by an auto-disconnect occurring while a mirror is running, the service is shutdown while a mirror is running, or if you stop a mirror manually.
- **Mirror complete**—This script starts when a mirror is completed. Because the mirror statistics may indicate a mirror is at 100% when it is actually still processing (for example, if files were added after the replication set size was calculated, if there are alternate data streams, and so on), the script will not start until all of the mirror data has been completely processed on the target.

1. Open the Replication Console.
2. Right-click the connection on the right pane of the Replication Console and select **Connection Manager**.
3. Select the **Scripts** tab.

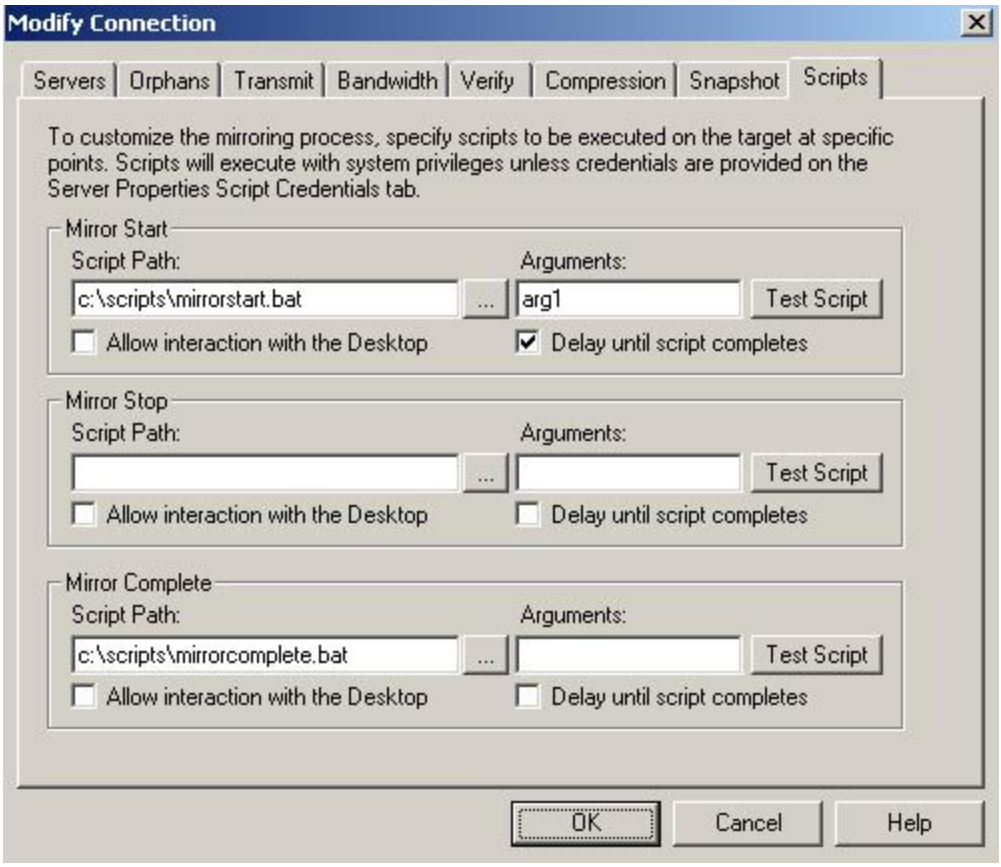

4. For each of the three predefined points in the mirroring process, specify the following information.
   - **Script path**—Specify the path and filename for each script.
   - **Arguments**—If needed, specify any arguments that are required to execute your script. The arguments must be valid for your script.
   - **Allow interaction with the Desktop**—Enable this option if you want the script processing to be displayed on the screen. Otherwise, the script will execute silently in the background.
   - **Delay until script completes**—Enable this option if you want to delay all Double-Take Availability processing until the script finishes processing.
   - **Test Script**—You can test your script manually by clicking **Test Script**. Your script will be executed if you test it. If the test is successful, the **Test Script** button will become disabled. If necessary, manually undo any changes that you do not want on your target after testing the script.
5. Click **OK** to save the settings.

Note: Mirror scripts are dependent on the target and target path location of a connection. Therefore, if you establish mirror scripts for one connection and then

establish additional connections to the same target using the same target path location, the mirror scripts will automatically be applied to those subsequent connections. If you select a different target path location, the mirror scripts will have to be reconfigured for the new connection(s).

Removing orphan files

An orphan file is a file that exists in the target's copy of the replication set data, but it does not exist in the source replication set data. An orphan file can be created when you delete a file contained in the source replication set while there is no Double-Take Availability connection. For example, if a connection was made and a mirror was completed and then the connection was stopped and a file was deleted on the source, an orphan file will exist on the target. Because the connection has been disconnected, the delete operation is not replicated to the target and the file is not deleted on the target. Additionally, orphan files may also exist if files were manually copied into or deleted from the location of the target's copy of the replication set data.

You can configure orphan files to be moved or deleted automatically during a mirror, verify, or restore, or you can move or delete orphan files manually at any time. You can move or delete all orphan files on the target or only those orphan files that are older than a specified period of time. The results of orphan processing are maintained in the Double-Take Availability log on the target, including the number of moved/deleted orphan files, the directories, and the number of bytes.

Note: Orphan file configuration is a per target option. All connections to the same target will have the same orphan file configuration.

The orphans feature does not move or delete alternate data streams. To do this, use a full mirror which will delete the additional stream(s) when the file is re-created.

If Double-Take Availability is configured to move orphan files, the Double-Take Availability log file will indicate that orphan files have been deleted even though they have actually been moved. This is a reporting issue only.

If delete orphans is enabled, directories and files that do not exist on the source and are excluded in the replication set using a wildcard rule will be removed from the target path. If you have data in your target path that does not exist on the source, do not use wildcard rules in your replication set. Manually select and deselect those files which should be included or excluded from your replication set.

---

1. Open the Replication Console.
2. If you want to remove orphan files manually, right-click an established connection and select **Remove Orphans**, **Start**.
3. If you want to stop the process after it has been started, right-click the connection and select **Remove Orphans**, **Stop**.

4. To configure orphan files for processing during a mirror, verify, or restore, use the following instructions.
    a. Right-click the connection on the right pane of the Replication Console and select **Connection Manager**.
    b. Select the **Orphans** tab.

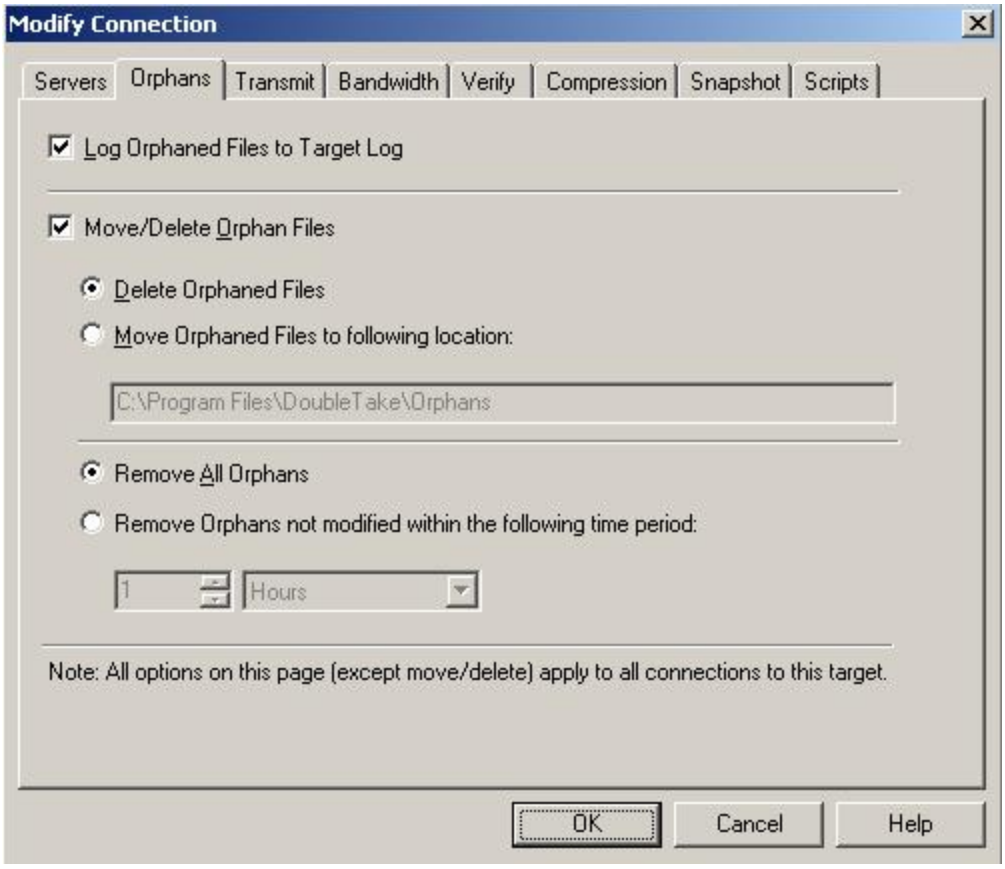

    c. Specify if you want to log the name of the orphan files to the Double-Take Availability log file on the target by marking **Log Orphaned Files to Target Log**.
    d. By default, the orphan files feature is disabled. To enable it, mark **Move/Delete Orphan Files**.
    e. Specify if you want to **Delete Orphaned Files** or **Move Orphaned Files** to a different location. If you select the move option, identify the location where these orphan files will be located.

Note: If you are moving files, make sure the directory you specify to move the files to is not included in the destination of the replication set data

so that the orphan files are only moved once.

f. Specify if you want to **Remove All Orphans** or **Remove Orphans not modified within the following time period**. If you select the time-based option, only orphans older than the time you specify will be removed.
g. Click **OK** to save the settings.

Replication

Replication is one of the key components of Double-Take Availability. This section contains the following replication topics.

- Replication capabilities—Review this list to learn what Double-Take Availability supports for replication.
- Replication sets—This section contains instructions for creating and using Double-Take Availability replication sets.
- Starting replication—Since replication is one of the key components of Double-Take Availability, this topic includes instructions for starting replication.
- Inserting tasks during replication—You can insert tasks to be processed inline with replication.

Replication capabilities

Double-Take Availability replicates file and directory data stored on any Windows file system (FAT, FAT32, NTFS4, and NTFS5). Replicated items also include Macintosh files, compressed files, NTFS attributes and ACLs (access control list), dynamic volumes, files with alternate data streams, sparse files, and encrypted files. Files can be replicated across mount points, even though mount points are not created on the target. Some reparse points are replicated, including CommVault Data Migrator and BridgeHead Software HT FileStore.

Double-Take Availability does not replicate items that are not stored on the file system, such as physical volume data and registry based data. Additionally, Double-Take Availability does not replicate NTFS extended attributes, registry hive files, Windows or any system or driver pagefile, system metadata files ($LogFile, $Mft, $BitMap, $Extend\\$UsnJrnl, $Extend\\$Quota, $Extend\\$ObjId, and $Extend\\$Reparse), hard links, or the Double-Take Availability disk-based queue logs. The only exception to these exclusions is for the full-server workloads. If you are protecting your system state and data using full-server protection, Double-Take Availability will automatically gather and replicate all necessary system state data, including files for the operating system and applications.

Note the following replication caveats.

1. If you have mixed file systems, keep in the mind the following.
    a. If, on your source, you have a FAT volume mounted on a directory which resides on an NTFS volume, these files will not be mirrored, regardless of the target file system. Replication will work correctly. To work around this issue, make sure both volumes are NTFS.
    b. If you are mirroring/replicating from an NTFS source to a FAT target, you may see additional error messages in your Double-Take Availability log file because the target file system cannot handle the NTFS attributes or file permissions. For example, if your replication set contains files with alternate data streams, you will see messages indicating that there are unfinished operations because the FAT file system cannot store the alternate data stream information.
    c. If you select a compressed file or folder from an NTFS partition and replicate it to a FAT target, the attributes are lost, but the data is maintained.
2. If any directory or file contained in your replication set specifically denies permission to the system account or the account running the Double-Take service, the attributes of the file on the target will not be updated because of the lack of access. This also includes denying permission to the Everyone group because this group contains the system account.

3. If you select a dynamic volume and you increase the size of the volume, the target must be able to compensate for an increase in the size of the dynamic volume.
4. If you select files with alternate data streams, keep in mind the following.
    a. Alternate data streams are not included in the replication set size calculation. Therefore, you may see the mirror process at 100% complete while mirroring continues.
    b. The number of files and directories reported to be mirrored will be incorrect. It will be off by the number of alternate streams contained in the files and directories because the alternate streams are not counted. This is a reporting issue only. The streams will be mirrored correctly.
    c. Use the checksum option when performing a difference mirror or verification to ensure that all alternate data streams are compared correctly.
    d. If your alternate streams are read-only, the times may be flagged as different if you are creating a verification report only. Initiating a remirror with the verification will correct this issue.
5. If you select encrypted files, keep in mind the following.
    a. Only the data, not the attributes or security/ownership, is replicated. However, the encryption key is included. This means that only the person who created the encrypted file on the source will have access to it on the target.
    b. Only data changes cause replication to occur; changing security/ownership or attributes does not.
    c. Replication will not occur until the Windows Cache Manager has released the file. This may take awhile, but replication will occur when Double-Take Availability can access the file.
    d. When remirroring, the entire file is transmitted every time, regardless of the remirror settings.
    e. Verification cannot check encrypted files because of the encryption. If remirror is selected, the entire encrypted file will be remirrored to the target. Independent of the remirror option, all encrypted files will be identified in the verification log.
    f. Empty encrypted files will be mirrored to the target, but if you copy or create an empty encrypted file within the replication set after mirroring is complete, the empty file will not be created on the target. As data is added to the empty file on the source, it will then be replicated to the target.
    g. When you are replicating encrypted files, a temporary file is created on both the source and target servers. The temporary file is automatically created in the same directory as the Double-Take Availability disk queues. If there is not enough room to create the temporary file, an out of disk space message will be logged. This message may be misleading and indicate that the drive

where the encrypted file is located is out of space, when it actually may be the location where the temporary file is trying to be created that is out of disk space.

6. If you are using mount points, keep in mind the following.
    a. By default, the mount point data will be stored in a directory on the target. You can create a mount point on the target to store the data or maintain the replicated data in a directory. If you use a directory, it must be able to handle the amount of data contained in the mount point.
    b. Recursive mount points are not supported. If you select data stored on a recursive mount point, mirroring will never finish.
7. Double-Take Availability supports transactional NTFS (TxF) write operations, with the exception of TxF SavePoints (intermediate rollback points).
    a. With transactional NTFS and Double-Take Availability mirroring, data that is in a pending transaction is in what is called a transacted view. If the pending transaction is committed, it is written to disk. If the pending transaction is aborted (rolled back), it is not written to disk.

       During a Double-Take Availability mirror, the transacted view of the data on the source is used. This means the data on the target will be the same as the transacted view of the data on the source. If there are pending transactions, the Double-Take Availability **Target Data Stat**e will indicate **Transactions Pending**. As the pending transactions are committed or aborted, Double-Take Availability mirrors any necessary changes to the target. Once all pending transactions are completed, the **Target Data State** will update to **OK**.

       If you see the pending transactions state, you can check the Double-Take Availability log file for a list of files with pending transactions. As transactions are committed or aborted, the list is updated until all transactions are complete, and the **Target Data State** is **OK**.

    b. During replication, transactional operations will be processed on the target identically as they are on the source. If a transaction is committed on the source, it will be committed on the target. If a transaction is aborted on the source, it will be aborted on the target.
    c. When failover occurs any pending transactions on the target will be aborted before the source identity is assigned to the target.
    d. Double-Take Availability restore functions as a mirror, except the roles of the source and target are reversed. The transacted view of the data on the target is restored to the source. As pending transactions are committed or aborted on the target, Double-Take Availability restores any necessary changes to

the source. Once all pending transactions are completed, the restoration is complete and replication will continue from the target to the source.

e. If you have restored your data before starting the failback process, make sure the restoration process does not have pending transactions and is complete before starting failback. If you are restoring your data after the failback the process has completed, users will not be accessing the data once failback occurs, so there are no opportunities for pending transactions.

8. Double-Take Availability supports Windows 2008 symbolic links and junction points. A symbolic link is a link (pointer) to a file. Junction points are also links, but to folders and volumes.
   a. If the link and the file/folder/volume are both in your source replication set, both the link and the file/folder/volume are mirrored and replicated to the target.
   b. If the link is in the source replication set, but the file/folder/volume it points to is not, only the link is mirrored and replicated to the target. The file/folder/volume that the link points to is not mirrored or replicated to the target. A message is logged to the Double-Take Availability log identifying this situation.
   c. If the file/folder/volume is in the source replication set, but the link pointing to it is not, only the file/folder/volume is mirrored and replicated to the target. The link pointing to the file/folder/volume is not mirrored or replicated to the target.
9. Short file names are not supported on FAT file systems.

Replication sets

A replication set defines the data on a source machine that Double-Take Availability protects. Replication sets are defined by volumes, directories, files, or wild card combinations. Creating multiple replication sets allows you to customize sets of data that need to be protected. When working with data workloads, you need to define the replication set data yourself. If you are protecting full-server, application, virtual, or cluster workloads, the replication set is automatically defined for you.

When a replication set is created, a series of rules are defined that identify the volumes, directories, files, and/or wild card combinations that will be replicated to the target. Each rule includes the following.

- **Path**—The path including volume, drive, directory, file, and/or wild card
- **Include**—If the specified path is to be included in the files sent to the target
- **Exclude**—If the specified path is to be excluded from the files sent to the target
- **Recursive**—If the rule should automatically be applied to the subdirectories of the specified path

For example, a replication set rule might be volume\directory* inc, rec

This specifies that all files contained in the volume\directory path are included in the replication set. Because recursion is set, all files and subdirectories under volume\directory are also included. A complete replication set becomes a list of replication set rules.

Replication sets offer flexibility tailoring Double-Take Availability to your environment. For example, multiple replication sets can be created and saved for a source to define a unique network configuration. There may be three replication sets - Critical Data, User Data, and Offsite Data. Critical Data could be configured to replicate, in real-time, to an onsite high-availability server. Offsite Data is replicated across a WAN and, therefore, is configured to queue changes until a sufficient amount of data is changed to justify transmission. At that point, the connection is made and stays active until all the data is transmitted. User Data is not replicated throughout the day, but a nightly changed file mirror copies only blocks of data that are different between the source and target server prior to a nightly tape backup operation being run on the target server. Each of these replication sets can be automated to transmit as needed, thus protecting your entire environment.

- Creating a replication set
- Creating or modifying replication rules manually
- Modifying a replication set
- Renaming and copying a replication set

- Calculating replication set size
- Deleting a replication set

Keep in mind the following notes when creating and working with replication sets and connections.

- **Limitations**
  - Replication set rules are limited in length meaning that the entire volume\directory\filename including slashes, spaces, periods, extensions, cannot exceed 259 characters.
  - Double-Take Availability can mirror, replicate, verify, and restore paths up to 32760 characters, although each individual component (file or directory name) is limited to 259 characters. Paths longer than 32760 characters will be skipped and logged.
  - Do not name replication sets or select a target location using illegal characters. Illegal characters include the following.
    - period .
    - question mark ?
    - forward or backward angle bracket < >
    - colon :
    - quotation mark "
    - forward or backward slash \ /
    - asterisk *
    - pipe or vertical bar |
- **Error checking and avoidance**
  - Do not connect more than one replication set to the same location on a target. You could overwrite or corrupt your data.
  - Replication sets contain error checking to avoid inadvertent overwrites of the replication set rules. When replication sets are modified, a generation number is associated with the modifications. The generation number is incremented anytime the modifications are saved, but the save is not allowed if there is a mismatch between the generation number on the source and the Replication Console. You will be notified that the replication set could not be saved. This error checking safeguards the replication set data in the event that more than one client machine is accessing the source's replication sets.
  - Double-Take Availability will not replicate the same data from two different replication sets on your source. The data will only be replicated from one of the replication sets. If you need to replicate the same data more than once, connect the same replication set to multiple targets.

- If you rename the root folder of a connected replication set, Double-Take Availability interprets this operation as a move from inside the replication set to outside the replication set. Therefore, since all of the files under that directory have been moved outside the replication set and are no longer a part of the replication set, those files will be deleted from the target copy of the replication set. This, in essence, will delete all of your replicated data from the target. If you have to rename the root directory of your replication set, make sure that the replication set is not connected.
- When creating replication sets, keep in mind that when recursive rules have the same type (include or exclude) and have the same root path, the top level recursive rule will take precedence over lower level non-recursive rules. For example, if you have /data included recursively and /data/old included nonrecursively, the top level rule, /data/, will take precedence and the rule /data/old will be discarded. If the rules are different types (for example, /data is included and /data/old is excluded), both rules will be applied as specified.

- **Including and excluding files**
  - Do not exclude Microsoft Office temporary files from your replication set. When a user opens a Microsoft Office file, a temporary copy of the file is opened. When the user closes the file, the temporary file is renamed to the original file and the original file is deleted. Double-Take Availability needs to replicate both the rename and the delete. If you have excluded the temporary files from your replication set, the rename operation will not be replicated, but the delete operation will be replicated. Therefore, you will have missing files on your target.
  - When Microsoft SQL Server databases are being replicated, you should always include the tempdb files, unless you can determine that they are not being used by any application. Some applications, such as PeopleSoft and BizTalk, write data to the tempdb file. You can, most likely, exclude temporary databases for other database applications, but you should consult the product documentation or other support resources before doing so.
  - Some applications create temporary files that are used to store information that may not be necessary to replicate. If user profiles and home directories are stored on a server and replicated, this could result in a significant amount of unnecessary data replication on large file servers. Additionally, the \Local Settings\Temporary Internet Files directory can easily reach a few thousand files and dozens of megabytes. When this is multiplied by a hundred users it can quickly add up to several gigabytes of data that do not need to be replicated.
  - Creating replication sets that only contain one file may cause unexpected results. If you need to replicate just one file, add a second file to the

replication set to ensure the data is replicated to the correct location. (The second file can be a zero byte file if desired.)

- **Backups**
  - Double-Take Availability does not replicate the last access time if it is the only thing that has changed. Therefore, if you are performing incremental or differential backups on your target machine, you need to make sure that your backup software is using an appropriate flag to identify what files have been updated since the last backup. You may want to use the last modified date on the file rather than the date of the last backup.
- **Virus protection**
  - Virus protection software on the target should not scan replicated data. If the data is protected on the source, operations that clean, delete, or quarantine infected files will be replicated to the target by Double-Take Availability. If the replicated data on the target must be scanned for viruses, configure the virus protection software on both the source and target to delete or quarantine infected files to a different directory that is not in the replication set. If the virus software denies access to the file because it is infected, Double-Take Availability will continually attempt to commit operations to that file until it is successful, and will not commit any other data until it can write to that file.

Creating a replication set

Before you can establish a connection, you must create a replication set.

1. From the Replication Console, highlight a source in the left pane of the Replication Console and select **Insert**, **Replication Set** from the menu bar. You can also right-click on the source name and select **New**, **Replication Set**.
2. A replication set icon appears in the left pane under the source. By default, it is named New Replication Set. Rename the newly inserted replication set with a unique name by typing over the default name and pressing **Enter**. This process is similar to naming a new folder in Windows Explorer.
3. Expand the tree under the replication set name to view the volume and directory tree for the source.

Note: The default number of files that are listed in the right pane of the Replication Console is 2500, but this is user configurable. A larger number of file listings allows you to see more files in the Replication Console, but results in a slower display rate. A smaller number of file listings displays faster, but may not show all files contained in the directory. To change the number of files displayed, select **File**, **Options** and adjust the **File Listings** slider bar to the desired number.

To hide offline files, such as those generated by snapshot applications, select **File**, **Options** and disable **Display Offline Files**. Offline files and folders are denoted by the arrow over the lower left corner of the folder or file icon.

---

4. Identify the data on the source that you want to protect by selecting volumes, drives, directories, and/or specific files.

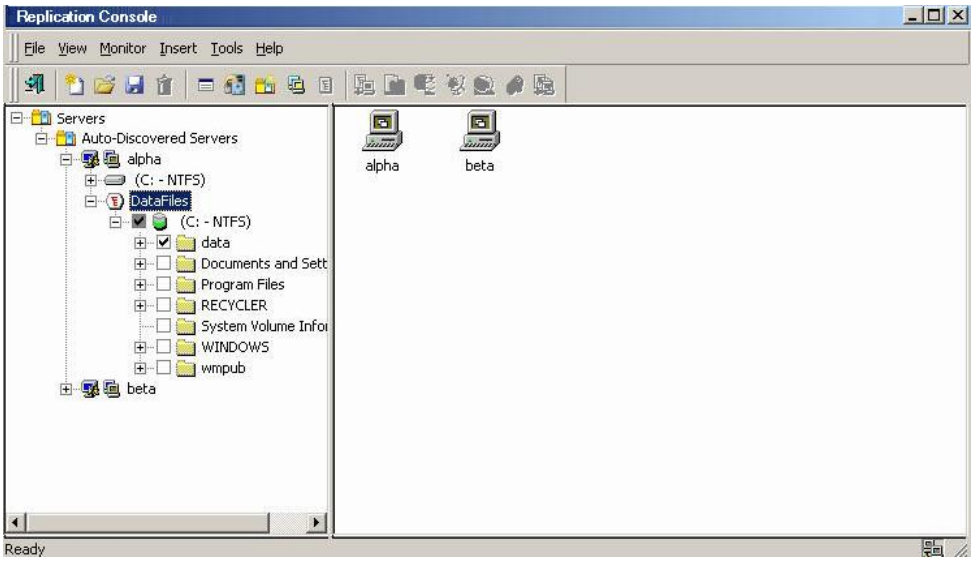

Note: Be sure and verify what files can be included by reviewing Replication capabilities.

5. After selecting the data for this replication set, right-click the new replication set icon and select **Save**. A saved replication set icon will change from red to black.

Creating or modifying replication rules manually

There may be times when you cannot browse for data when creating a replication set. For example, you can create a replication set rule for a directory or file that does not exist. Since you cannot browse for the location, you have to create replication set rule manually. At other times, the data you want to replicate cannot be easily selected from the Replication Console. For example, you may want to select all .db files from a specific volume or directory. This task may be easier to complete by creating the replication set rule manually. Use the following instructions to create or modify a replication set rule manually.

1. Open the Replication Console.
2. If you do not have a replication set created, you need to create one. Highlight a source in the left pane of the Replication Console and select **Insert**, **Replication Set** from the menu bar. You can also right-click on the source name and select **New**, **Replication Set**. A replication set icon appears in the left pane under the source. By default, it is named New Replication Set. Rename the newly inserted replication set with a unique name by typing over the default name and pressing **Enter**. This process is similar to naming a new folder in Windows Explorer.
3. Right-click on the replication set icon and select **Properties**. The Replication Set Properties dialog box appears and lists any existing rules. The existing rules may have been entered manually or selected by browsing the source. Each rule will display the attributes associated it.

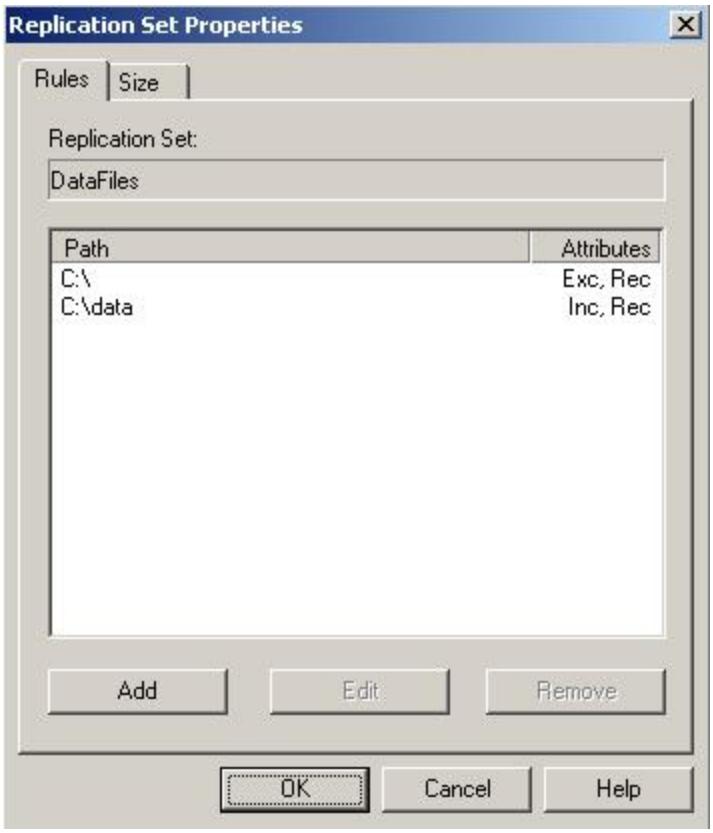

- **Inc**—Include indicates that the specified path is to be included in the files sent to the target
- **Exc**—Exclude indicates that the specified path is not to be included in the files sent to the target
- **Rec**—Recursion indicates the rule should automatically be applied to the subdirectories of the specified path. If you do not select this option, the rule will not be applied to subdirectories.

4. From the Replication Set Properties dialog box, click **Add**.
5. Specify a path, wild card, or specific file name. Select the **Include, Exclude**, and/or **Recurse sub-directories** attributes to be applied to this rule and click **OK**.
6. Repeat steps 3 and 4 to add additional replication set rules.
7. If you need to edit an existing rule, highlight it and click **Edit**.
8. If you need to remove a rule, highlight it and click **Remove**.
9. After the replication set rules have been defined, exit the Replication Set Properties dialog box by clicking **OK**. Notice the replication set icon has changed from black to red, indicating changes to the replication set rules. If you click **Cancel**, your changes will not be reflected in the current replication set.
10. Right-click the replication set icon and select **Save**. A saved replication set icon will change from red to black.

Modifying a replication set

Double-Take Availability allows you to make modifications to a replication set when you want to change the data you wish to protect. This allows you to add, remove, or modify any replication set rules without having to create a new replication set.

1. Open the Replication Console.
2. In the left pane, highlight the replication set you want to modify and expand the volume and directory levels as needed.
3. Modify the items by marking or clearing the volume, drive, directory, or file check boxes. Notice the replication set icon has changed from black to red, indicating changes to the replication set rules.
4. After updating the rules for this replication set, right-click the replication set icon and select **Save**. A saved replication set icon will change from red to black.

Note: If you save changes to a connected replication set, it is recommended that you perform a mirror to guarantee data integrity between the source and target machines. A dialog box will appear instructing you to disconnect and reconnect the replication set and perform a difference mirror.

If your source is a cluster, you must make the same modifications to the replication set on the non-owning nodes. The replication set rules must be identical on all nodes for Double-Take Availability to function properly on the cluster.

Renaming and copying a replication set

To rename or copy a replication set, open the Replication Console. Click once on a highlighted replication set name to edit the field. Specify a unique name and press **Enter**. The process is similar to renaming a folder in Windows Explorer. If the original replication set has not been saved (red icon), the new name replaces the original name. If the original replication set is saved (black icon), the new name creates a copy of the original replication set.

Note: If you save changes to a connected replication set, it is recommended that you perform a mirror to guarantee data integrity between the source and target machines. A dialog box will appear instructing you to disconnect and reconnect the replication set and perform a difference mirror.

If your source is a cluster, you must make the same modifications to the replication set on the non-owning nodes. The replication set rules must be identical on all nodes for Double-Take Availability to function properly on the cluster.

Calculating replication set size

While Double-Take Availability is mirroring, the right pane of the Replication Console displays statistics to keep you informed of its progress. If the size of the replication set is determined before the mirror is started, Double-Take Availability can display the percentage of the replication set that has been mirrored in the **Mirror Status** column. If the size was not calculated prior to starting the mirror, the column displays **Mirroring**.

1. Open the Replication Console.
2. Right-click on the replication set icon and select **Properties**. The Replication Set Properties dialog box appears.
3. Select the **Size** tab.

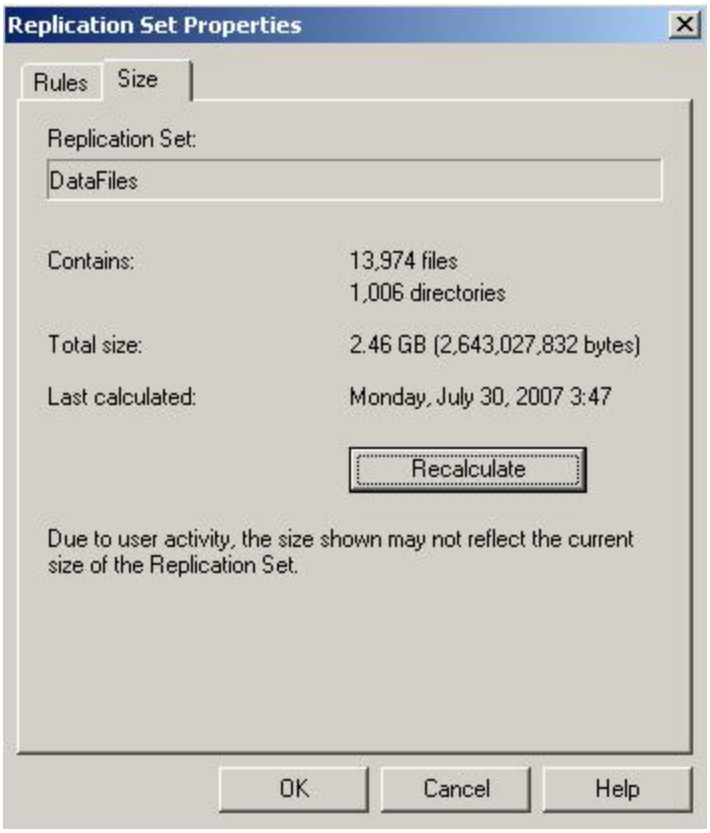

4. If the replication set size has never been determined, click **Calculate**. If the replication set has previously been determined, the button will be labeled **Recalculate**. Depending on user activity, the size shown may not accurately reflect the current size of the replication set. If changes are occurring to files in the replication set while the calculation is being made, the actual size may differ slightly. The amount of data is determined at the exact time the calculation is made.
5. Click **OK** to return to the Replication Console.

Note: You can also configure the replication set calculation when establishing a connection through the Connection Manager by selecting **Calculate Replication Set size** on connection on the **Mirroring** tab.

If your replication set contains a large number of files, for example, ten thousand or more, you may want to disable the calculation of the replication set size so that data will start being mirrored sooner. If calculation is enabled, the source calculates the file size before it starts mirroring. This can take a significant amount of time depending on the number of files and system performance. Disabling calculation will result in the mirror status not showing the percentage complete or the number of bytes remaining to be mirrored.

Deleting a replication set

You can only delete a replication set if it is not currently connected. If the replication set is connected, you must disconnect the connection and then delete the replication set.

To delete a replication set, open the Replication Console. Right-click the replication set icon and select **Delete**. Additionally, you can highlight the replication set and press the **Delete** key on the keyboard.

Starting replication

Starting replication when establishing a connection is the default and recommended configuration. If replication is not started, data is not added to the queue on the source, and source/target data integrity is not guaranteed.

To start replication, open the Replication Console. Right-click the connection on the right pane of the Replication Console and select **Replication**, **Start**. After starting replication, you should perform a remirror to guarantee the source and target data are identical.

Inserting tasks during replication

Task command processing is a Double-Take Availability feature that allows you to insert and run tasks at various points during the replication of data. Because the tasks are user-defined, you can achieve a wide variety of goals with this feature. For example, you might insert a task to create a snapshot or run a backup on the target after a certain segment of data from the source has been applied on the target. This allows you to coordinate a point-in-time backup with real-time replication.

Task command processing can be enabled from the Replication Console, but it can only be initiated through the scripting language. See the *Scripting Guide* for more information.

To enable task command processing open the Replication Console. Right-click a server in the left pane of the Replication Console, select **Properties**, select the **Setup** tab, and select **Enable Task Command Processing**.

Note: If you disable this option on a source server, you can still submit tasks to be processed on a target, although task command processing must be enabled on the target.

Verification

Verification is the process of confirming that the data on the target is identical to the data on the source. Verification creates a log file detailing what was verified as well as which files are not synchronized. If the data is not the same, Double-Take Availability can automatically initiate a remirror. The remirror ensures data integrity between the source and target.

Note: Because of the way the Windows Cache Manager handles memory, machines that are doing minimal or light processing may have file operations that remain in the cache until additional operations flush them out. This may make Double-Take Availability files on the target appear as if they are not synchronized. When the Windows Cache Manager releases the operations in the cache on the source, the files will be updated on the target.

- Verifying manually—You can verify your data at any time manually.
- Verifying on a schedule—You can schedule verification tasks for periodic intervals.
- Configuring the verification log—You can configure how the verification information is logged.
- Verify applications on the target—You can verify that you Exchange or SQL application database on the target is viable for failover.

Verifying manually

A manual verification can be run anytime a mirror is not in progress.

1. Open the Replication Console.
2. Right-click the connection on the right pane of the Replication Console and select **Verify**.
3. Select the verification options that you would like to perform.

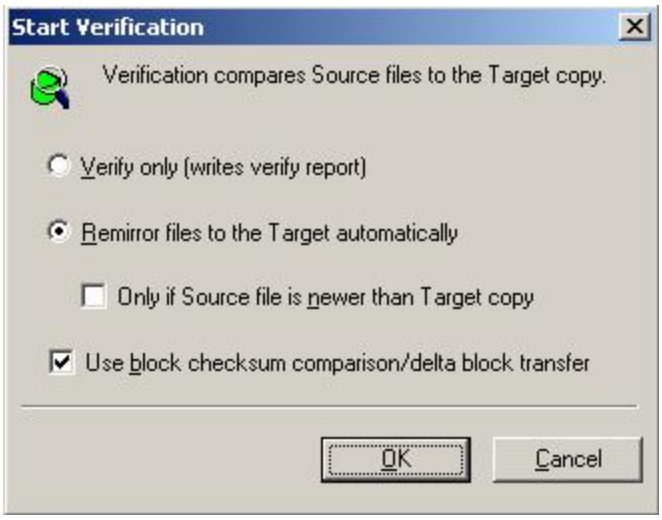

- **Verify only**—This option verifies the data and generates a verification log, but it does not remirror any files that are different on the source and target.
- **Remirror files to the Target automatically**—This option verifies the data, generates a verification log, and remirrors to the target any files that are different on the source.
- **Only if the Source file is newer than Target copy**—If you are remirroring your files, you can specify that only files that are newer on the source than the target be remirrored.

Note: If you are using a database application, do not use the newer option unless you know for certain you need it. With database applications, it is critical that all files, not just some of them that might be newer, get mirrored.

- **Use block checksum comparison/delta block transfer**—Specify if you want the verification process to use a block checksum comparison to determine which blocks are different. If this option is enabled, only those

blocks (not the entire files) that are different will be identified in the log and/or remirrored to the target.

Note: Database applications may update files without changing the date, time, or file size. Therefore, if you are using database applications, you should use the block checksum comparison to ensure proper verification and remirroring.

---

4. Click **OK** to start the verification.

Verifying on a schedule

Verification can be scheduled to occur automatically at periodic intervals.

1. Open the Replication Console.
2. Right-click the connection on the right pane of the Replication Console and select **Connection Manager**.
3. Select the **Verify** tab.

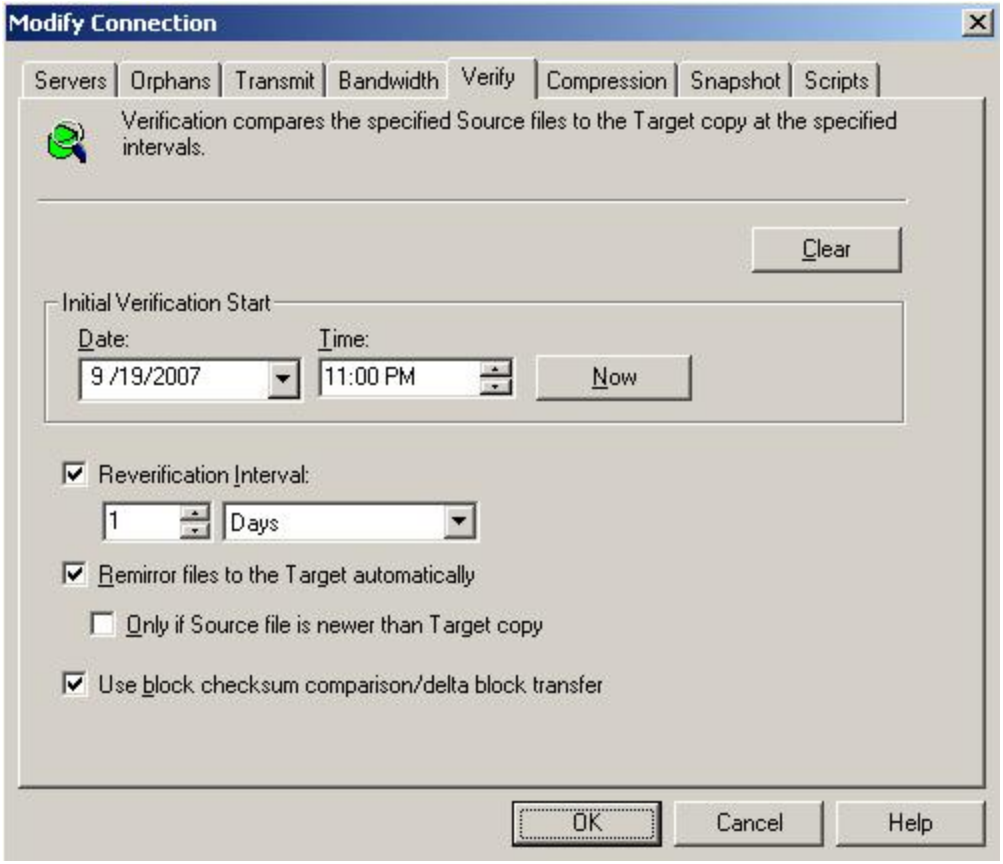

4. Specify when you want to start the initial verification. Select the immediate date and time by clicking **Now**, or enter a specific **Date** and **Time**. The down arrow next to **Date** displays a calendar allowing easy selection of any date. **Time** is formatted for any AM or PM time.
5. Mark the **Reverification Interval** check box to repeat the verification process at the specified interval. Specify an amount of time and choose minutes, hours, or days.
6. Select if you want to **Remirror files to the Target automatically**. When enabled, Double-Take Availability will verify the data, generate a verification log, and remirror to the target any files that are different on the source. If disabled, Double-Take Availability will verify the data and generate a verification log, but no files will be remirrored to the target.

7. If you are remirroring your files, you can specify **Only if Source file is newer than Target copy** so that only files that are newer on the source than on the target are remirrored.

Note: If you are using a database application, do not use the newer option unless you know for certain you need it. With database applications, it is critical that all files, not just some of them that might be newer, get mirrored.

---

8. Specify if you want the verification process to **Use block checksum comparison/delta block transfer** to determine which blocks are different. If this option is enabled, only those blocks (not the entire files) that are different will be identified in the log and/or remirrored to the target.

Note: Database applications may update files without changing the date, time, or file size. Therefore, if you are using database applications, you should use the block checksum comparison to ensure proper verification and remirroring.

---

9. Click **OK** to save the settings.

Configuring the verification log

A verification log is created on the source during the verification process. The log identifies what is verified as well as which files are not synchronized.

1. Open the Replication Console.
2. Right-click the source server on the left pane of the Replication Console and select **Properties**.
3. Select the **Logging** tab.

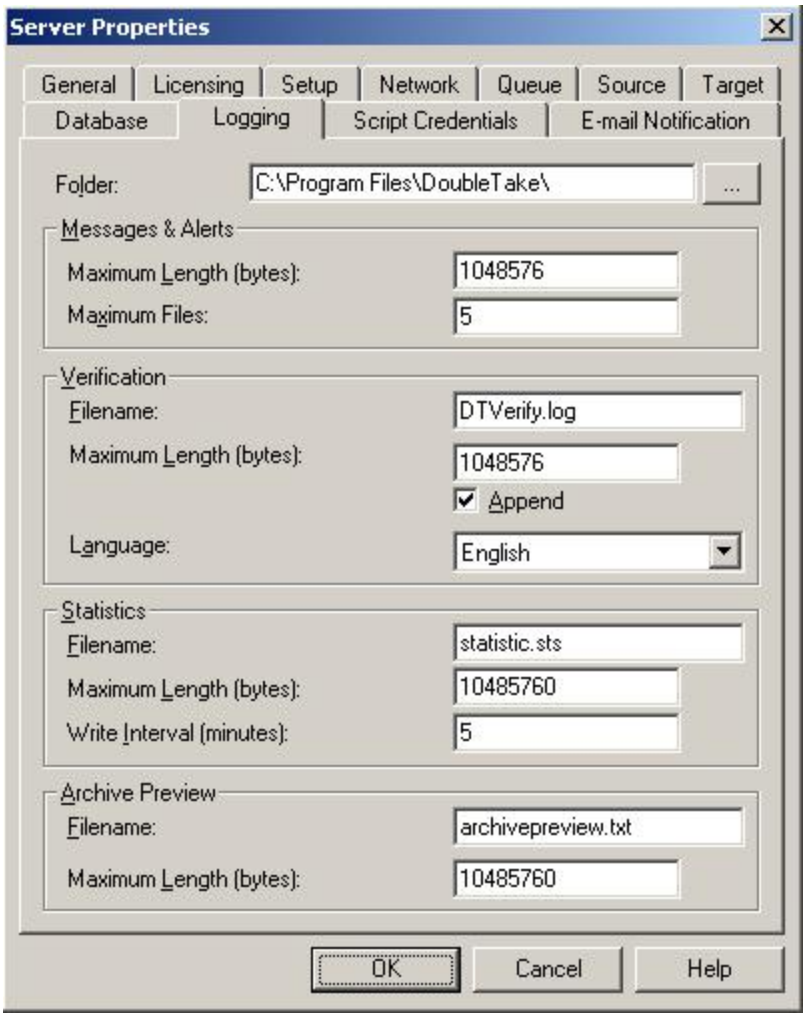

4. At the top of the window, **Folder** identifies the location where the log files identified on this tab are stored. By default, the log files are stored in the same directory as the Double-Take Availability program files.
5. Under the Verification section, **Filename** contains the base log file name for the verification process. The replication set name will be prefixed to the base log file

name. For example, since the default is DTVerify.log, the verification log for the replication set called UserData would be UserData DTVerify.log.

6. Specify the **Maximum Length** of the log file. The default is 1048576 bytes (1 MB). When the log file reaches this limit, no additional data will be logged.
7. By default, the log is appended to itself each time a verification process is completed. Clear the **Append** check box if you do not want to append to the previous log file.

Note: Changes made to the verification log in the **Server Properties**, **Logging** tab will apply to all connections from the current source machine.

---

8. Specify the **Language** of the log file. Currently, English is the only available language.
9. Click **OK** to save the settings.

In the log file, each verification process is delineated by beginning and end markers. A list of files that are different on the source and target is provided as well cumulative totals for the verification process. The information provided for each file is the state of its synchronization between the source and the target at the time the file is verified. If the remirror option is selected so that files that are different are remirrored, the data in the verify log reflects the state of the file before it is remirrored, and does not report the state of the file after it is remirrored. If a file is reported as different, review the output for the file to determine what is different.

Sample verification log

```
--- VERIFICATION OF CONNECTION 2 (Sales data for alpha --> 206.31.65.40 :
1100) ---
Start Time: 1/24/2010 12:15:20 PM for connection 2 (Sales data for alpha -->
206.31.65.40 : 1100)
File:       beta\users\bob\budget.xls DIFFERENT ON TARGET
     Source Attributes: Timestamp = 1/17/2010 8:21:36 PM Size = 1272 Mask =
[0x20]
     Target Attributes: Timestamp = 1/17/2010 8:21:36 PM Size = 1272 Mask =
[0x20]
Security descriptors are different.
      0 BYTES OUT OF SYNC
File:       beta\users\bill\timesheet.xls DIFFERENT ON TARGET
     Source Attributes: Timestamp = 1/17/2010 8:21:37 PM Size = 1272 Mask =
[0x20]
     Target Attributes: Timestamp = 1/17/2010 8:21:37 PM Size = 1272 Mask =
[0x23]
      0 BYTES OUT OF SYNC
```

```
File:       beta\users\vincent\training.doc DIFFERENT ON TARGET
     Source Attributes: Timestamp = 1/12/2010 3:28:20 PM Size = 17 Mask =
[0x20]
     Target Attributes: Timestamp = 1/20/2010 5:05:26 PM Size = 2 Mask =
[0x20]
      17 BYTES OUT OF SYNC
Completion Time: 1/24/2010 12:37:44 PM for connection 2 (Sales data for
alpha -->
206.31.65.40 : 1100)
Elapsed Time (seconds): 1320.256470
Total Directories Compared: 657
Total Directories Missing: 0
Total Directories Remirrored: 0
Total Files Compared: 120978
Total Files Missing: 0
Total Files Different: 3
Total Files Encrypted: 0
Total Files Remirrored: 1
Total Bytes Skipped: 0
Total Bytes Compared: 18527203678
Total Bytes Missing: 0
Total Bytes Different: 17
Total Bytes Remirrored: 17
Related links and directory attributes have been adjusted.
----- END OF VERIFICATION -----
```

- **Timestamp**—The last modified date and time of the file
- **Size**—The size, in bytes, of the file
- **Mask**—The attributes associated with the file. See further details below.
- **Security descriptors**—The NTFS file permissions of the file. If the file permissions are different, the message "Security descriptors are different" will be logged. If the file permissions are the same, nothing will be logged.
- **Bytes out of sync**—The number of bytes that are not synchronized between the file on the source and the file on the target. If the data in the file is identical, the message "0 BYTES OUT OF SYNC" will be logged. If the file is different, the message will indicate how many bytes were different. This message does not indicate that the file was remirrored during the verify.

The mask must be converted in order to determine what attributes are assigned to a file. The mask is a hexadecimal number corresponding to a binary number that indicates what the attributes are. Using the following steps, you can determine how the mask corresponds to the attributes of a file.

1. Each mask begins with 0x. Identify the hexadecimal number after the constant 0x. For example, if the mask is 0x23, then the hexadecimal number you are interested in is 23. The hexadecimal number may be up to four digits.
2. Convert the hexadecimal number to its 16-digit binary equivalent. You can use the Windows calculator for this conversion.
    a. Select **Start**, **Programs**, **Accessories**, **Calculator**.
    b. Switch to scientific view, if it is not already in that view, by selecting **View**, **Scientific**.
    c. Select **Hex**.
    d. Enter the hexadecimal number, for example 23, as specified in your verification log.
    e. Select **Bin** and the hexadecimal number will change to the binary equivalent.
    f. Pad the beginning of the binary equivalent with zeroes (0) so that the number is 16 digits long. For example, hexadecimal number 23 converts to 100011, so the 16-digit binary equivalent would be 0000000000100011.
3. Determine what number (0 or 1) appears in each position of the binary number. Because binary numbers count from right to left, start with position 1 on the right.
    - 1—Read only
    - 2—Hidden
    - 3—None
    - 4—System
    - 5—Directory
    - 6—Archive
    - 7—Encrypted
    - 8—Normal
    - 9—Temporary
    - 10—Sparse file
    - 11—Reparse point
    - 12—Compressed
    - 13—Offline
    - 14—Not content indexed
    - 15—None
    - 16—None
4. Using the list above, identify those attributes that are enabled by those positions equal to one (1). The positions equal to zero (0) are disabled and that attribute does not apply. So hexadecimal number 23, which converted to

0000000000100011, indicates read only, hidden, and archive. Another example might be mask 0x827 which converted to binary is 0000100000100111. Positions 1-3, 6, and 12 are all enabled which indicates the file is read only, hidden, archive, and compressed.

Note: Files that were replicated with the **Replicate NT Security by Name** feature enabled, will be identified as different in the log file because of the local name attribute. The files will be the same.

Verify applications on the target

The application verification process confirms that an Exchange or SQL application database on the target is viable for failover.

Note: The application verification process is not available for Exchange or SQL in a cluster environment.

The application verification process can only be performed on active connections that are in a good state and are not failed over.

---

1. Verify that your current volumes have adequate space to contain snapshots of your target. These snapshots will be used to revert the target back to its pre-test state after you have completed your application verification.
2. Disable target path blocking for your application protection connection.
3. Verify that your Double-Take disk queue is not located on a volume that you will be reverting (a volume with application data). Also verify that the queues have adequate space to handle the data changes that will be queued while you are testing your application.
4. If you are verifying Exchange and you are running the Application Manager from a machine other than the source or target, you must install Exchange System Manager on that machine.
5. From the Application Manager, select **Actions**, **Verify Target Data**.
6. There are four main sections of the Database Verification window.
   - **Status**—The top of the window displays the overall status of the application verification. You can click on the status description for more detailed information.
   - **Results**—Initially, the state of the application verification is unknown, as indicated by the question mark icon. When the databases or stores have been successfully mounted, the states will update to green.
   - **History**—This area shows the sequence of verification events.
   - **Options**—You can select which services are tested and specify scripts to run during the test.
   - **Start core services only**—Only the core application services will be started when the test is performed.
   - **Start selected services**—All of the services you configured for application failover will be started when the test is performed. Use this option if you have an application add-on such as BlackBerry.

- **Script to run after target is online**—Specify a script, located in the Double-Take Availability installation folder on the target, to run after the target application is online. A sample script to move users, that you can modify to fit your environment, is available in the \Samples sub-directory of your Double-Take Availability installation.
- **Script to run before restoring normal protection**—Specify a script, located in the Double-Take Availability installation folder on the target, to run before stopping the application on the target. A sample script to move users, that you can modify to fit your environment, is available in the \Samples sub-directory of your Double-Take Availability installation.

7. Click **Test** to begin the application verification. The **Status**, **Results**, and **History** will update during the test. When the **Status** is **Target online**, the verification is complete. You can perform any custom testing at this time.
8. If you have any problems during the test, click **Undo** to revert the target to its pre-test state. All snapshots and items created for the verification test will be removed. If you do not have Volume Shadow Copy installed, the snapshots used during the test will not be removed.
9. When you have completed any custom testing, click **Continue** to revert back to the pre-test state. All snapshots and items created for the verification test will be removed. If you do not have Volume Shadow Copy installed, the snapshots used during the test will not be removed.

Verifying applications on the target from the command line

If desired, you can verify your application on the target by using the TDV utility from the command line.

Command

TDV

Description

Utility that confirms that the application database on the target is viable for failover

Syntax

TDV /APPTYPE <SQL | EXCHANGE> /DNSDOMAIN <*domain_name*> /SRCNAME <*source_name*> /TARNAME <*target_name*> /MODE <INSTANCE|DATABASE> [/PORT <*port_number*>] [/USERNAME <*user_name*>] [/PASSWORD <*password*>] /SVC <APP|ALL> [/ADDONSVC <*service1,service2, ..*>] [/SETPASSWORD <*username*> <*password*>] [/GETPASSWORD] [/SCRIPTPOST <*post_online*>] [/SCRIPTPRE <*pre_restore*>] /SRCEXCHVER <2003|2007> /TAREXCHVER <2003|2007> /SRCVER <2000|2005|2008|MSDE|EXPR> /TARVER <2000|2005|2008|MSDE|EXPR> [/INTERACTIVE] [/HELP]

Options

- APPTYPE—Specify the keyword SQL or EXCHANGE to indicate the application being verified on the target
- DNSDOMAIN <*domain_name*>—Fully-qualified name of the domain
- SRCNAME <*source_name*>—Name of the source server
- TARNAME <*target_name*>—Name of the target server
- MODE—Specify the keyword INSTANCE or DATABASE to indicate the type of SQL database being verified
- PORT <*port_number*>—The source and target port number. This port number must be the same on both servers.
- USERNAME <*user_name*>—Name of the user login account
- PASSWORD <*password*>—The password associated with specified user name
- SVC—Specify the keyword APP or ALL to indicate if only the core application services will be started when the test is performed or if all of

the services you configured for application failover will be started when the test is performed. Use the ALL option if you have an application add-on such as BlackBerry.

- ADDONSVC <***service1,service2, ..***>—Specify any additional services to run when the test is performed
- SETPASSWORD <***username***> <***password***>—Stores the specified user name and password in an encrypted file for later use
- GETPASSWORD—Retrieves the user name and password previously stored with the /SETPASSWORD option
- SCRIPTPOST <***post_online***>—A script, located in the Double-Take Availability installation folder on the target, to run after the target application is online. If there are spaces in the path and/or filename, enclose the path in quotation marks.
- SCRIPTPRE <***pre_restore***>—A script, located in the Double-Take Availability installation folder on the target, to run before stopping the application on the target. If there are spaces in the path and/or filename, enclose the path in quotation marks.
- SRCEXCHVER <2003|2007>—Specify the keyword 2003 or 2007 to indicate the version of Exchange running on the source
- TAREXCHVER <2003|2007>—Specify the keyword 2003 or 2007 to indicate the version of Exchange running on the target
- SRCVER <2000|2005|2008|MSDE|EXPR>—Specify the keyword 2000, 2005, 2008, MSDE, or EXPR to indicate the version of SQL running on the source
- TARVER <2000|2005|2008|MSDE|EXPR>—Specify the keyword 2000, 2005, 2008, MSDE, or EXPR to indicate the version of SQL running on the target
- INTERACTIVE—Runs the TDV utility in interactive mode. You will be prompted to continue before the test is started and after the test is complete.
- HELP—Displays command syntax help

Examples

- TDV /apptype EXCH /dnsdomain corp.greek.com /srcname alpha /tarname beta /username administrator /password password /svc app /interactive
- TDV /apptype SQL /srcname alpha /tarname beta /mode instance /username administrator /password password /svc all

Notes

- The application verification process can only be performed on active connections that are in a good state and are not failed over.
- If you deselected a SQL instance when configuring SQL protection, all of the named instances on the target will be listed, including any you deselected.
- If Application Manager is open when the TDV utility is run and becomes unresponsive, close and reopen Application Manager.
- Any status other than online or offline will be reported as unknown.

Data transmission

Double-Take Availability data is continuously transmitted to the target machine. Although the data may be queued if the network or target machine is slow, the default transmission setting is to transmit the data as soon as possible. You can modify the transmission to suit your environment.

- Stopping, starting, pausing, and resuming transmission—You can maintain the source/target connection, but still control the transmission of data across the network by using the manual transmission controls. If transmission is paused, the data is queued on the source until you manually restart the transmission.
- Scheduling data transmission—You can set event driven or scheduling criteria to determine when data is transmitted. Data is queued on the source until the event or schedule is met. Also, transmission can be stopped by using these criteria. Scheduled transmission options can be toggled on and off, allowing you to enable them only when you need to use them.
- Limiting transmission bandwidth—You can specify bandwidth limitations to restrict the amount of network bandwidth used for Double-Take Availability data transmissions. Data is queued on the source until bandwidth is available. Bandwidth limitations can be full-time or scheduled.
- Compressing data for transmission—You can compress data to reduce the amount of bandwidth needed to transmit Double-Take Availability data.

Stopping, starting, pausing, and resuming transmission

To start, pause, or resume the transmission of data from the source to the target, open the Replication Console. Right-click an established connection, select **Transmit** and the appropriate transmission control.

Scheduling data transmission

Using the Connection Manager **Transmit** tab, you can set start and stop criteria along with a schedule window.

Note: Double-Take Availability checks the schedule once every second, and if a user-defined criteria is met, transmission will start or stop, depending on the option specified.

Any replication sets from a source connected to the same IP address on a target will share the same scheduled transmission configuration.

---

1. Open the Replication Console.
2. Right-click the connection on the right pane of the Replication Console and select **Connection Manager**.
3. Select the **Transmit** tab. The **Transmit** tab contains four limit types: **Bandwidth**, **Start**, **Stop**, and **Window**. The transmission options for each limit type are displayed by highlighting a selection in the **Limit Type** box.

   At the top of the **Transmit** tab dialog box, the **Enable Transmission Limiting** check box allows you to turn the transmission options on or off. You can enable the transmission options by marking the **Enable Transmission Limiting** check box when you want the options to be applied, but you can disable the transmission options, without losing the settings, by clearing that check box.

   Also at the top of the **Transmit** tab dialog box, the **Clear All** button, when selected, will remove all transmission limitations that have been set under any of the limit types. The **Clear** button will clear the settings only for the **Limit Type** selected.
4. When you schedule transmission start criteria, transmission will start when the criteria is met and will continue until the queue is empty or a transmission stop criteria is met. Select the **Start option** in the **Limit Type** box.

   Define the start options for Double-Take Availability transmission by using any combination of the following options.

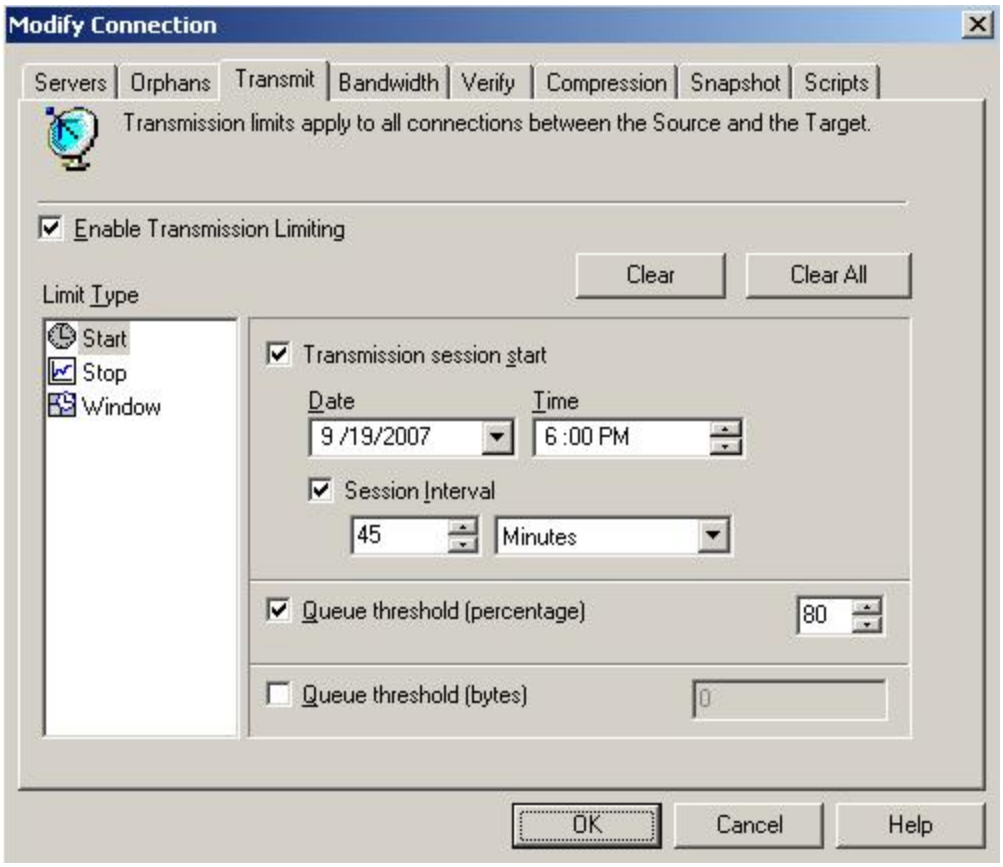

- **Transmission session start**—This option establishes a date and time of the day to begin transmitting data. For example, you may want to specify a transmission time that corresponds to a low bandwidth usage time. Once started, Double-Take Availability will continue to transmit data until the queue is empty or until another limitation stops the transmission. Specify a **Date** and **Time** to start transmitting data. The down arrow next to the date field displays a calendar allowing easy selection of any date. The time field is formatted for any AM or PM time.
- **Session Interval**—This option begins transmitting Double-Take Availability data at specified intervals of time. This option is used in conjunction with **Transmission session start**. For example, if the **Session Interval** is set to repeat transmission every 30 minutes and the **Transmission session start** is set to begin transmitting at 10 PM, if the queue is emptied at 10:20 the transmission will stop. The start criteria is again met at 10:30 and Double-Take Availability will begin transmitting any new data in the queue. Specify an interval for additional transmissions by indicating a length of time and choosing minutes, hours, or days.
- **Queue Threshold (percentage) and Queue threshold (bytes)**—If the allocated amount of queue disk space is in use, Double-Take Availability

cannot continue to queue data causing an auto-disconnect and the potential for loss of data. To avoid using the entire queue, you can configure Double-Take Availability to begin transmitting data to the target when the queue reaches a certain point. This point can be defined as a percentage of the disk queue that must be in use or the number of bytes in the disk queue. For example, if you specify 40%, when 40% of the queue is in use, Double-Take Availability initiates the transmission process and sends the data in the queue to the target machine. The transmission stops when the queue is empty or a Double-Take Availability stop transmission criteria is met. Or you might set a queue threshold of 500 MB. Double-Take Availability will wait until there is 500 MB of data in the queue and then begin transmitting the data. Like other start criteria, Double-Take Availability continues transmitting until the queue is empty or a Double-Take Availability stop criteria is met. Specify a percentage of the disk queue and system memory that must be in use to initiate the transmission process, and/or specify the number of bytes that must be in the source queue and system memory to initiate the transmission process.

Note: A **Transmission Session Start** setting will override any other start criteria. For example, if you set the **Transmission Session Start** and the **Queue Threshold**, transmission will not start until you reach the indicated start time.

---

5. Schedule any desired stop criteria to stop transmission after a transmission start criteria has initiated the transmission. If you do not establish a stop criteria, transmission will end when the queue is empty. Select the **Stop** option in the **Limit Type** box.

   Define the stop options to stop Double-Take Availability transmissions by using either or both of the following options.

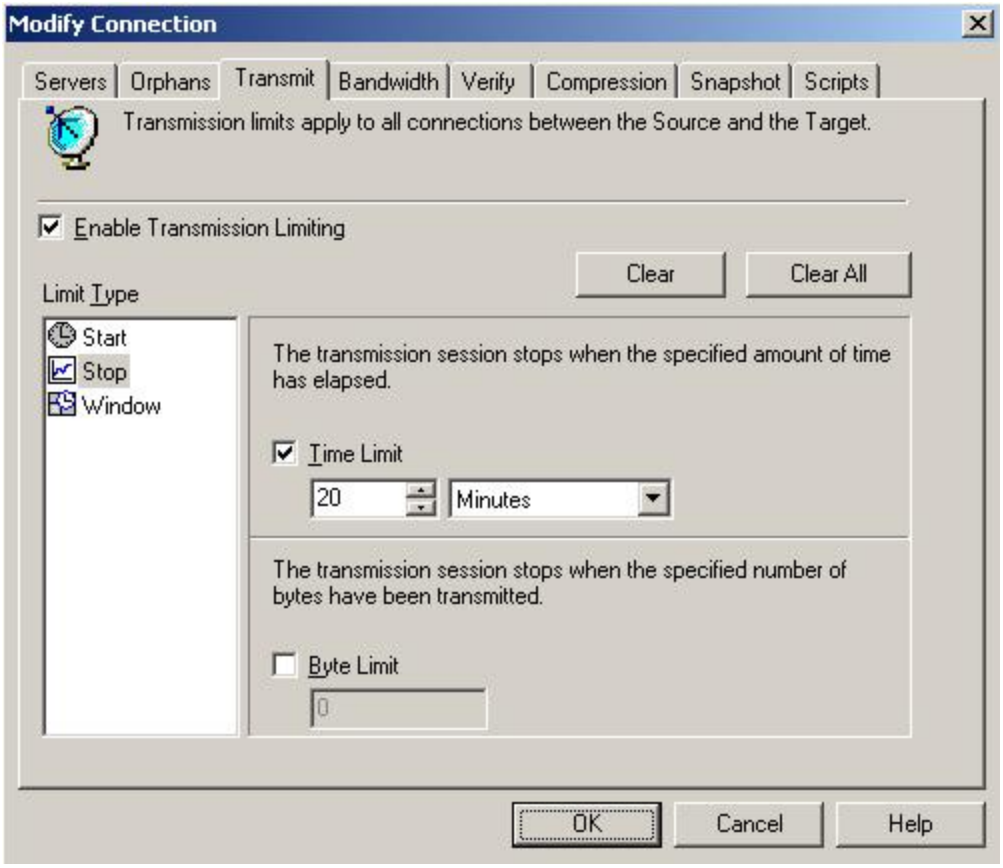

- **Time Limit**—The time limit specifies the maximum length of time for each transmission period. Any data that is not sent during the specified time limit remains on the source queue. When used in conjunction with the session interval start option, you can explicitly define how often data is transmitted and how long each transmission lasts. Specify the maximum length of time that Double-Take Availability can continue transmitting by indicating a length of time and choosing minutes, hours, or days.
- **Byte Limit**—The byte limit specifies the maximum number of bytes that can be sent before ending the transmission session. When the byte limit is met, Double-Take Availability will automatically stop transmitting data to the target. Any data that still remains waits in the source queue until the transmission is restarted. When used in conjunction with a session start option, you can explicitly define how much data is being sent at a given time. Specify the maximum number of bytes that can be sent before ending the Double-Take Availability transmission.

Note: The transmission start and stop criteria should be used in conjunction with each other. For example, if you set the **Queue Threshold** equal to 10 MB and the **Byte Limit** equal to 10 MB, a network connection

will be established when there is 10 MB of data in the queue. The data will be transmitted and when the 10 MB **Byte Limit** is reached, the network connection closes. This is useful in configurations where metered charges are based on connection time.

6. Schedule a transmission window to establish a period of availability for all Double-Take Availability transmissions. You can specify a begin and end time for all Double-Take Availability transmissions. When a transmission window is in effect, all other start and stop criteria are bound by this window. This means that Double-Take Availability will never transmit data outside of an established window, regardless of other transmission settings. For example, if you set a window of availability from 9 p.m. to 4 a.m. and a start option to initiate transmission at 5 a.m., the window option will override the start option and no data will be sent at 5 a.m. Select the **Window** option in the **Limit Type** box.

   Note: Setting a transmission window by itself is not sufficient to start a transmission. You still need to set a start criteria within the window.

   Define a window to control Double-Take Availability transmissions by enabling the feature and then specifying both window options.

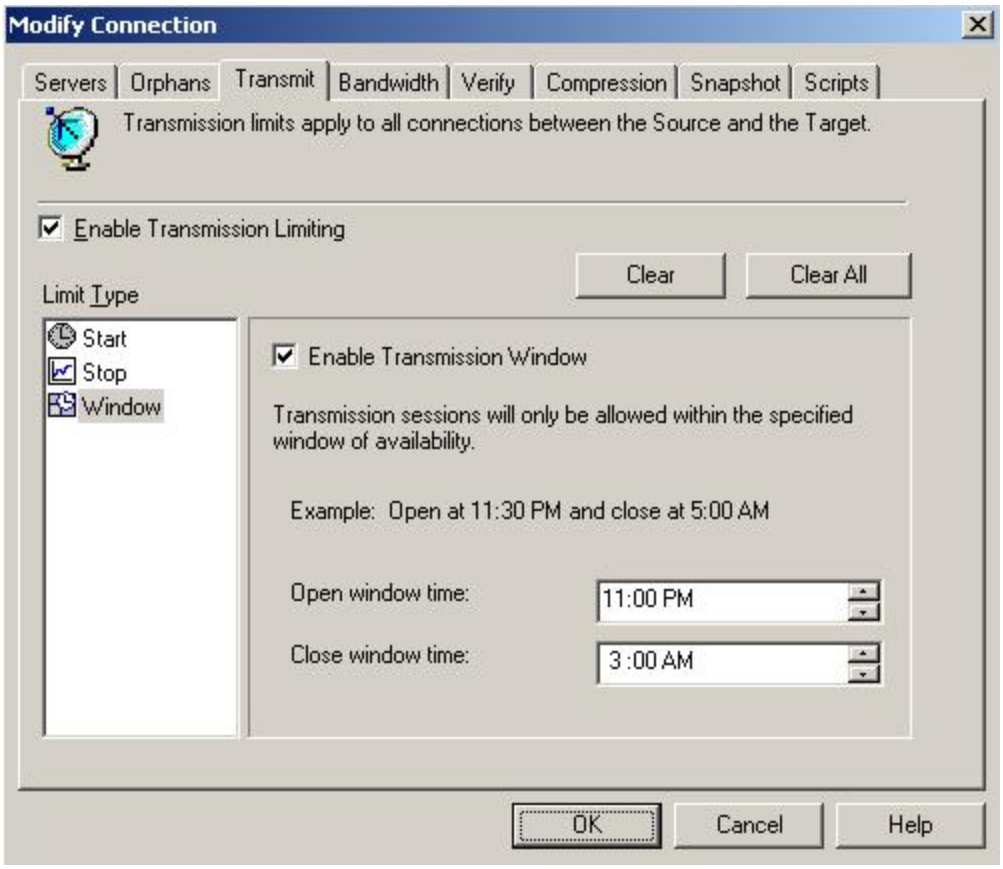

- **Enable Transmission Window**—This option specifies whether a transmission window is in use.
- **Open window time**—Specifies the time, formatted for AM or PM, when the transmission window will open, allowing transmission to begin.
- **Close window time**—Specifies the time, formatted for AM or PM, when the transmission window will close, stopping all transmission.

7. Click **OK** to save the settings.

Limiting transmission bandwidth

Bandwidth limitations are available to restrict the amount of network bandwidth used for Double-Take Availabilitydata transmissions. The network administrator specifies a percentage of bandwidth that is available or an absolute bandwidth limit for Double-Take Availability transmissions and Double-Take Availability never exceeds that allotted amount. The bandwidth not in use by Double-Take Availability is available for all other network traffic. You can schedule when you want bandwidth limiting to occur.

Note: Any replication sets from a source connected to the same IP address on a target will share the same bandwidth limitation configuration.

You will not be able to set a limit lower than 100% of a 28.8 Kbps connection speed. A setting this low would cause abnormal behavior between Double-Take Availability servers because of the lengthy delay between commands and responses transmitted between the two servers.

You cannot set a bandwidth limit of zero (0). If you need to stop data transmission completely, use the stop criteria on the Connection Manager **Transmit** tab.

---

1. Open the Replication Console.
2. Right-click the connection on the right pane of the Replication Console and select **Connection Manager**.
3. Select the **Bandwidth** tab.

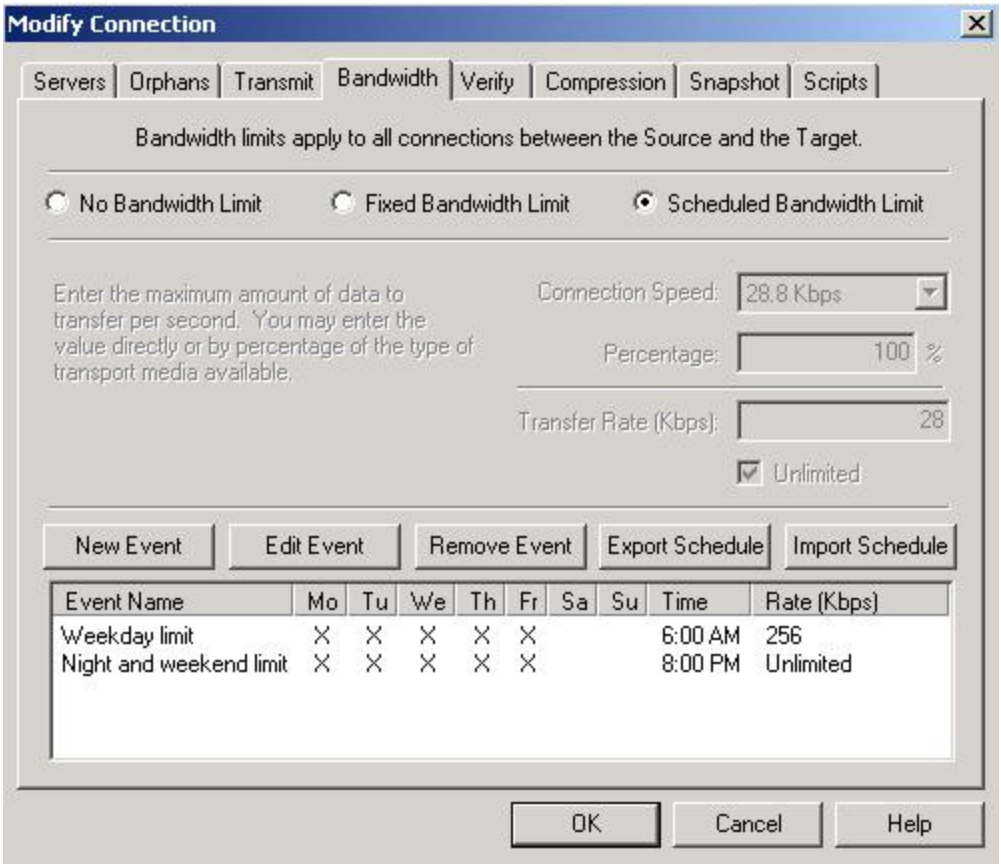

4. You have three bandwidth choices.
    - **No Bandwidth Limit**—Data will be transmitted at all times using all available bandwidth.
    - **Fixed Bandwidth Limit**—Data will be transmitted at all times according to the user-specified bandwidth configuration.
    - **Scheduled Bandwidth Limit**—Data will be transmitted according to the user-specified schedule and the user-specified bandwidth configuration.
5. If you want to transmit data at all times using all of the available bandwidth, select **No Bandwidth Limit**.
6. If you want to transmit data at all times using limited bandwidth, select **Fixed Bandwidth Limit**.
    a. By default, the **Unlimited** checkbox is enabled. This configuration is identical to selecting **No Bandwidth Limit**. If you want to limit your bandwidth usage, clear this checkbox.
    b. To limit the bandwidth usage, enter the maximum amount of data you want to transfer per second. You can indicate it by specifying your **Connection Speed** and the **Percentage** of the bandwidth that you want to use or by entering the **Transfer Rate** value directly.

7. If you want to transmit data according to a schedule using limited bandwidth, select **Scheduled Bandwidth Limit**.
    a. Click **New Event** to create a bandwidth schedule event.
        1. Specify a name for the bandwidth schedule event.
        2. Select the day(s) of the week that you want this event to be initiated on.
        3. Specify the time when you want this event to start.
        4. Specify the bandwidth limitation. By default, the **Unlimited** checkbox is enabled. This setting will not limit the transfer of Double-Take Availability data during the day and time specified. If you want to limit your bandwidth usage, clear this checkbox. To limit the bandwidth usage during the day and time specified, specify the maximum amount of data that you want to transfer per second. You can indicate it by specifying your **Connection Speed** and the **Percentage** of the bandwidth that you want to use or by entering the **Transfer Rate** value directly.
    b. Repeat each part of the new event creation to establish a comprehensive bandwidth schedule. For example, if you want to limit Double-Take Availability bandwidth usage weekdays (Monday - Friday) from 6:00 AM to 8:00 PM and have unlimited bandwidth usage outside of those times, you would create two events. The first event would be for Monday through Friday, would start at 6:00 AM, and would limit the bandwidth to the desired rate. The second event would be Monday through Friday, would start at 8:00 PM, and would not limit the bandwidth. (The Unlimited checkbox would be enabled.) In this schedule, on Monday at 6:00 AM, the schedule limit would be applied, limiting the transfer of data. At 8:00 PM that night, the schedule limit would be applied again, unlimiting the transfer of data. At 6:00 AM Tuesday morning, the schedule would be applied again. This would continue until the Friday 8:00 PM schedule is applied. That event would remain in effect until Monday morning at 6:00 AM.
    c. Highlight an existing schedule event and click **Edit Event** to make modifications to the event. The same steps used to create the event can be used to edit the event.
    d. Highlight an existing schedule event and click **Remove Event** to delete the event that is no longer needed.
    e. You can export the entire schedule for use on another Double-Take Availability server by clicking **Export Schedule**. Specify a file name and location, and an XML file with the schedule information will be saved.
    f. You can import an exported schedule by clicking **Import Schedule**. Locate the XML file and click OK. Any existing schedule will be overwritten by the imported schedule.

8. Click **OK** to save the settings.

Note: You can establish a bandwidth schedule and then disable or override it by selecting **No Bandwidth Limit** or **Fixed Bandwidth Limit**. The schedule criteria will be saved and will not be reactivated until you reselect **Scheduled Bandwidth Limit**.

You can modify the bandwidth limits applied to a connection that is already established by right-clicking on the connection and selecting **Set Bandwidth**. A modified version of the Connection Manager Bandwidth tab will allow you to select a different bandwidth limitation. You cannot create a schedule from this dialog box. An existing schedule must already exist in the Connection Manager.

Compressing data for transmission

To help reduce the amount of bandwidth needed to transmit Double-Take Availability data, compression allows you to compress data prior to transmitting it across the network. In a WAN environment this provides optimal use of your network resources. If compression is enabled, the data is compressed before it is transmitted from the source. When the target receives the compressed data, it decompresses it and then writes it to disk. On a default Double-Take Availability installation, compression is disabled.

Note: Any replication sets from a source connected to the same IP address on a target will share the same compression configuration.

Keep in mind that the process of compressing data impacts processor usage on the source. 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. Use the following guidelines to determine whether you should enable compression:

- If data is being queued on the source at any time, consider enabling compression.
- If the server CPU utilization is averaging over 85%, be cautious about enabling compression.
- The higher the level of compression, the higher the CPU utilization will be.
- Do not enable compression if most of the data is inherently compressed. Many image (.jpg, .gif) and media (.wmv, .mp3, .mpg) files, for example, are already compressed. Some images files, such as .bmp and .tif, are uncompressed, so enabling compression would be beneficial for those types.
- Compression may improve performance even in high-bandwidth environments.
- Do not enable compression in conjunction with a WAN Accelerator. Use one or the other to compress Double-Take Availability data.

Use the following instructions to configure data compression.

1. Right-click the connection on the right pane of the Replication Console and select **Connection Manager**.
2. Select the **Compression** tab.

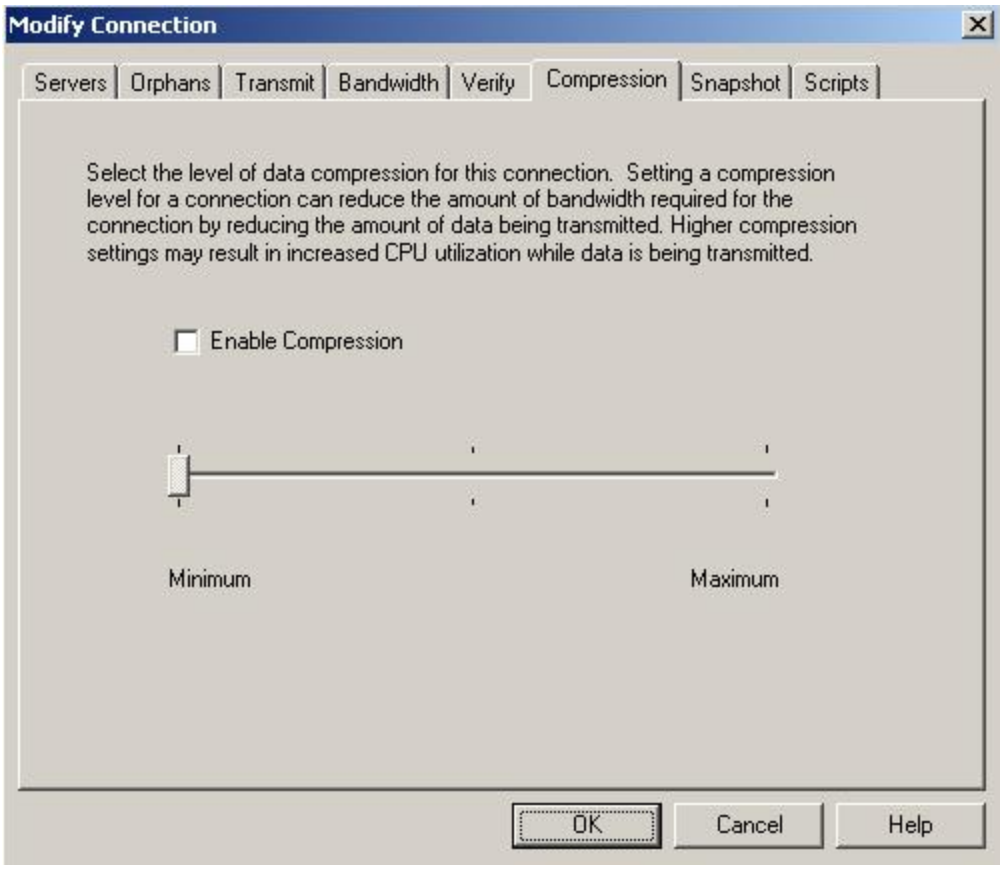

3. By default, compression is disabled. To enable it, select **Enable Compression**.
4. Depending on the compression algorithms available for your operating system, you may see a slider bar indicating different compression levels. Set the level from minimum to maximum compression to suit your needs.
5. Click **OK** to save the settings.

Snapshots

A snapshot is an image of data taken at a single point in time. Snapshots allow you to view files and folders as they existed at points of time in the past, so you can, for example, recover files that were accidentally deleted or overwritten. You could also compare a current revision of a file with an older revision. Double-Take Availability utilizes snapshot functionality by allowing you to create snapshots of the replicated data stored on the Double-Take Availability target.

Double-Take Availability snapshot functionality ensures that you will always have usable data on the target. For example, if your source server becomes infected with a virus, you can revert to a previous snapshot of the data on the target that was created prior to the virus infection. If you know the data on your target is good data, in a usable state, it will minimize application downtime in the event of a source failure. For example, if the source failed and the data on the target is not good due to an incomplete mirror, you can revert to a good snapshot on the target before failover. Snapshots also allow you to retrieve files that a user may have deleted.

Double-Take Availability uses the Microsoft Volume Shadow Copy service to create snapshots. To access this functionality, your target must be running Windows 2003 Service Pack 1 or later. Your servers must also be using the NTFS file system. Snapshots are taken at the volume level, corresponding to the target volumes contained in your replication set. For example, if your replication set contains d:\data and e:\files, the snapshot will contain all of the data on both the d: and e: volumes. If your replication set only includes d:\data (e:\files exists but is not included in the replication set), the snapshot will only contain the d: volume.

Sometimes taking a snapshot may not be possible. For example, there may not be enough disk space to create and store the snapshot, or maybe the target is too low on memory. If a snapshot fails, an Event message and a Double-Take Availability log message are both created and logged.

There are limitations imposed by Microsoft Volume Shadow Copy that impact Double-Take Availability snapshots. For example, Double-Take Availability maintains only 64 snapshots because Volume Shadow Copy only maintains 64 snapshots. If 64 snapshots exist and another one is taken, the oldest snapshots are deleted to make room for the new one. Another example is that Double-Take Availability snapshots must be created within one minute because Volume Shadow Copy snapshots must be created within one minute. If it takes longer than one minute to create the snapshot, the snapshot will be considered a failure. Additionally, Volume Shadow Copy will not revert snapshots of a volume with operating system files, therefore Double-Take Availability is also unable to revert a volume with operating system files. You must also keep in mind that if you are using extended functionality provided by Volume Shadow Copy, you need to be aware

of the impacts that functionality may have on Double-Take Availability. For example, if you change the location where the shadow copies are stored and an error occurs, it may appear to be a Double-Take Availability error when it is in fact a Volume Shadow Copy error. Be sure and review any events created by the VolSnap driver and check your Volume Shadow Copy documentation for details.

- Snapshots for data workloads
- Managing full-server and application snapshots

Snapshots for data workloads

- Snapshot states explains the various states of data workload snapshots.
- Automatic snapshots explains when automatic snapshots are taken.
- Scheduling snapshots contains instructions for scheduling periodic snapshots.
- Taking snapshots manually contains instructions for taking manual snapshots.

Snapshot states

A snapshot may not necessarily be useful if the data on the target is in a bad state. You only want snapshots of data that is in a good state. Therefore, you need to understand when the data is in a good or bad state.

Mirror started

- **State**—Bad
- **Description**—Mirroring has started, but is not complete. The data on the source and target will not be synchronized until the mirror is complete.
- **Automatic action taken for scheduled and automatic snapshots**—Scheduled and automatic snapshots will be delayed until the mirror is complete before taking a snapshot.
- **User interaction required for manual snapshots**—Wait until the mirror is complete and the data is in a good state, then take a manual snapshot.

Mirror stopped

- **State**—Bad
- **Description**—Mirroring has stopped without completing. The data on the source and target will not be synchronized until the mirror is complete.
- **Automatic action taken for scheduled and automatic snapshots**—Scheduled and automatic snapshots will be delayed until the mirror has been restarted and is complete before taking a snapshot.
- **User interaction required for manual snapshots**—Restart the mirror, wait until it is complete and the data is in a good state, and then take a manual snapshot.

Mirror complete

- **State**—Good
- **Description**—Because the mirror is complete, the data on the source and target is synchronized. Double-Take Availability will take a snapshot while the data is in a good state.
- **Automatic action taken for scheduled and automatic snapshots**—Scheduled and automatic snapshots will occur normally.
- **User interaction required for manual snapshots**—Manual snapshots can be taken normally.

Write operation retried

- **State**—Good
- **Description**—An operation cannot be written to the hard drive on the target. For example, the file could be in use by another application on the target.
- **Automatic action taken for scheduled and automatic snapshots**—Scheduled and automatic snapshots will occur normally, although the operation that is being retried will not be included in the snapshot.
- **User interaction required for manual snapshots**—Manual snapshots can be taken normally, although the operation that is being retried will not be included in the snapshot.

Write operation dropped

- **State**—Bad
- **Description**—An operation could not be written to the hard drive on the target, even after multiple retries. For example, the file could be in use by another application on the target.
- **Automatic action taken for scheduled and automatic snapshots**—An automatic snapshot will be taken just prior to the operation being dropped. Scheduled snapshots will be delayed until the target data is back in a good state.
- **User interaction required for manual snapshots**—Start a mirror, wait until it is complete and the data is in a good state, and then take a manual snapshot.

Write operation succeeded

- **State**—Good
- **Description**—An operation that was retrying on the target has been successfully written to the hard drive.
- **Automatic action taken for scheduled and automatic snapshots**—Scheduled and automatic snapshots will occur normally.
- **User interaction required for manual snapshots**—Manual snapshots can be taken normally.

Target restarted with connection persistence

- **State**—Good
- **Description**—The target service was able to persist connection information prior to restarting.

- **Automatic action taken for scheduled and automatic snapshots**—Scheduled and automatic snapshots will occur normally.
- **User interaction required for manual snapshots**—Manual snapshots can be taken normally.

Target restarted without connection persistence

- **State**—Bad
- **Description**—The target service has been restarted and was unable to persist connection information, therefore, operations that were in the queue have been lost.
- **Automatic action taken for scheduled and automatic snapshots**—An automatic snapshot will be taken after the target restarts, if the target data was in a good state prior to the target restart and the connection is configured to auto-remirror on auto-reconnect. Scheduled snapshots will be delayed until the target data is back in a good state.
- **User interaction required for manual snapshots**—Start a mirror, wait until it is complete and the data is in a good state, and then take a manual snapshot.

Restore required

- **State**—Good or bad
- **Description**—The data on the target no longer matches the data on the source because of a failover. This does not necessarily mean that the data on the target is bad.
- **Automatic action taken for scheduled and automatic snapshots**—Scheduled and automatic snapshots will be delayed until a restore is completed or the Restore Required state is overruled by a mirror. Once the restoration or mirror is complete, automatic and scheduled snapshots will occur normally.
- **User interaction required for manual snapshots**—Restore the target data back to the source or overrule the Restore Required state by performing a mirror. Once the restoration or mirror is complete, manual snapshots can be taken normally.

Snapshot reverted

- **State**—Good or bad
- **Description**—The data on the target no longer matches the data on the source because a snapshot has been applied on the target. This does not necessarily mean that the data on the target is bad.

- **Automatic action taken for scheduled and automatic snapshots**—Scheduled and automatic snapshots will be delayed until a restore is completed or the Snapshot Reverted state is overruled by a mirror. Once the restoration or mirror is complete, automatic and scheduled snapshots will occur normally.
- **User interaction required for manual snapshots**—Restore the target data back to the source or overrule the Snapshot Reverted state by performing a mirror. Once the restoration or mirror is complete, manual snapshots can be taken normally.

Restore complete

- **State**—Good
- **Description**—Because the restoration is complete, the data on the source and target is synchronized.
- **Automatic action taken for scheduled and automatic snapshots**—Scheduled and automatic snapshots will occur normally.
- **User interaction required for manual snapshots**—Manual snapshots can be taken normally.

To be completely assured that your data on the target is good, automatic and scheduled snapshots only occur when the data is in a good Double-Take Availability state. However, manual snapshots can be taken during any state. There are instances when you may want to take a manual snapshot, even if the target data is in a bad state. For example, if you drop an operation, that does not necessarily mean your data on the target is corrupt or the target would be unable to stand in for the source in the event of a failure. A snapshot of a bad state may be useful and usable, depending on your environment. If your source is a file server and an operation has been dropped, it is just one user file that is out-of-date. All of the remaining target files are intact and can be accessed in the event of a failure.

However, if your source is an application server and an operation has been dropped, that one file could cause the application not to start on the target in the event of a failure. In these cases, manual snapshots of a bad state depend on the context of your environment.

Note: Because the driver for Volume Shadow Copy is started before the driver for Double-Take Availability, if you revert any files on the source that are contained in your replication set, Double-Take Availability will not be aware of the revert and, therefore, the file change will not be replicated to the target. The file change will be mirrored to the target during the next mirroring process.

Automatic snapshots

When Double-Take Availability transitions from a good state to a bad state, it will automatically attempt to take a snapshot of the data before it leaves the good state and enters the bad state. For example, if your data is in a good state and you start a mirror, before the mirror is started, Double-Take Availability will automatically take a snapshot of the target. In the event the mirror fails to complete, you will have a snapshot of the data on the target when it was in its last good state.

Only one automatic snapshot per connection is maintained on the target. When an automatic snapshot is taken, it replaces any previous automatic snapshots.

Scheduling snapshots

You can schedule snapshots of your target data to fit your environment and needs. If the target data is in a bad state at the time of a scheduled snapshot, the snapshot will be delayed until the data on the target reaches a good state. If multiple scheduled snapshots are delayed, only one snapshot will be taken when the data reaches a good state.

1. Open the Replication Console.
2. Right-click the connection on the right pane of the Replication Console and select Connection Manager.
3. Select the **Snapshot** tab. You will not see the **Snapshots** tab if your server does not meet the snapshot requirements.

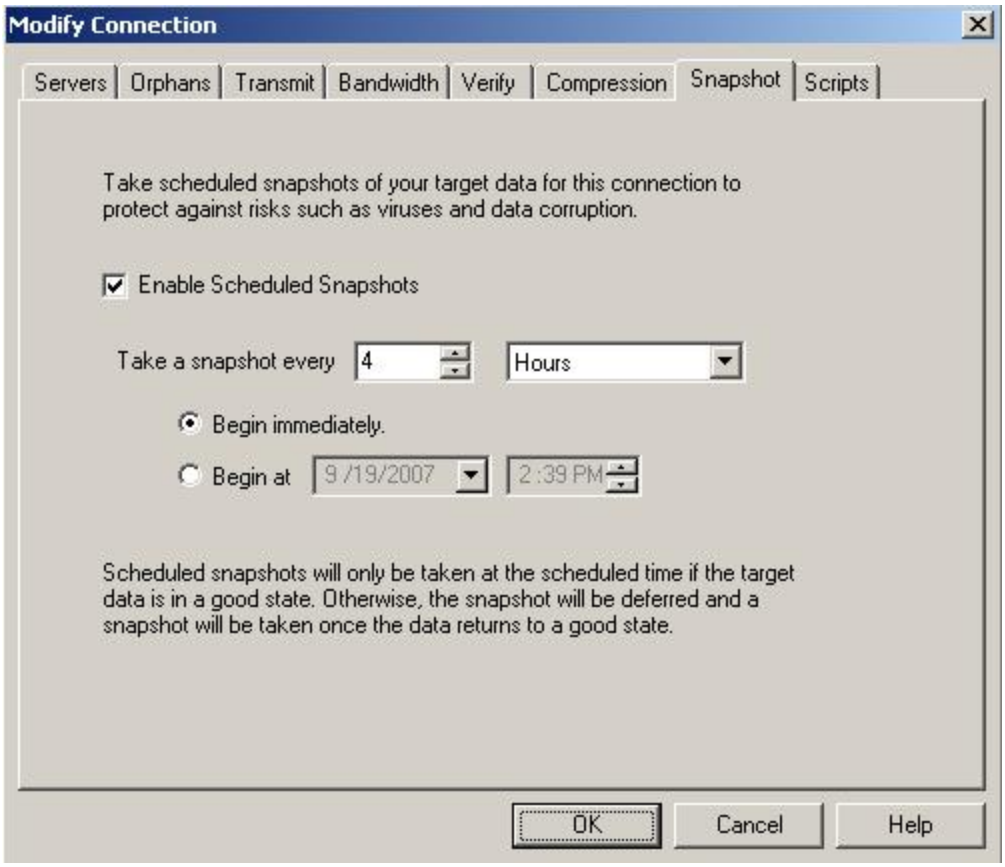

4. Select **Enable Scheduled Snapshots**.
5. Specify when snapshots should be taken, indicating minutes, hours, or days. The snapshots should be at least 15 minutes apart. By default, a snapshot is taken every hour.

6. Specify if you want the snapshots to start immediately. Otherwise, enter a date and time for when the snapshot schedule will begin.
7. Click **OK** to save the settings.

Taking snapshots manually

You can manually take a snapshot of the data on the target at any time. If an automatic or scheduled snapshot is currently in progress, Double-Take Availability will wait until that one is finished before taking the manual snapshot. Open the Replication Console, right-click the connection, and select **Snapshot Now**.

Managing full-server and application snapshots

By default, snapshots are enabled for full-server and application protections. Also by default, a snapshot is taken every 60 minutes. This may lead to numerous snapshots on the target that you may want to manage. You can do that from the Full-Server Failover Manger by selecting **Actions**, **Snapshot Manager** or from the Application Manager by selecting **Tools**, **Manage Snapshots**. (These options are only available when a source and target are selected and protection is enabled.)

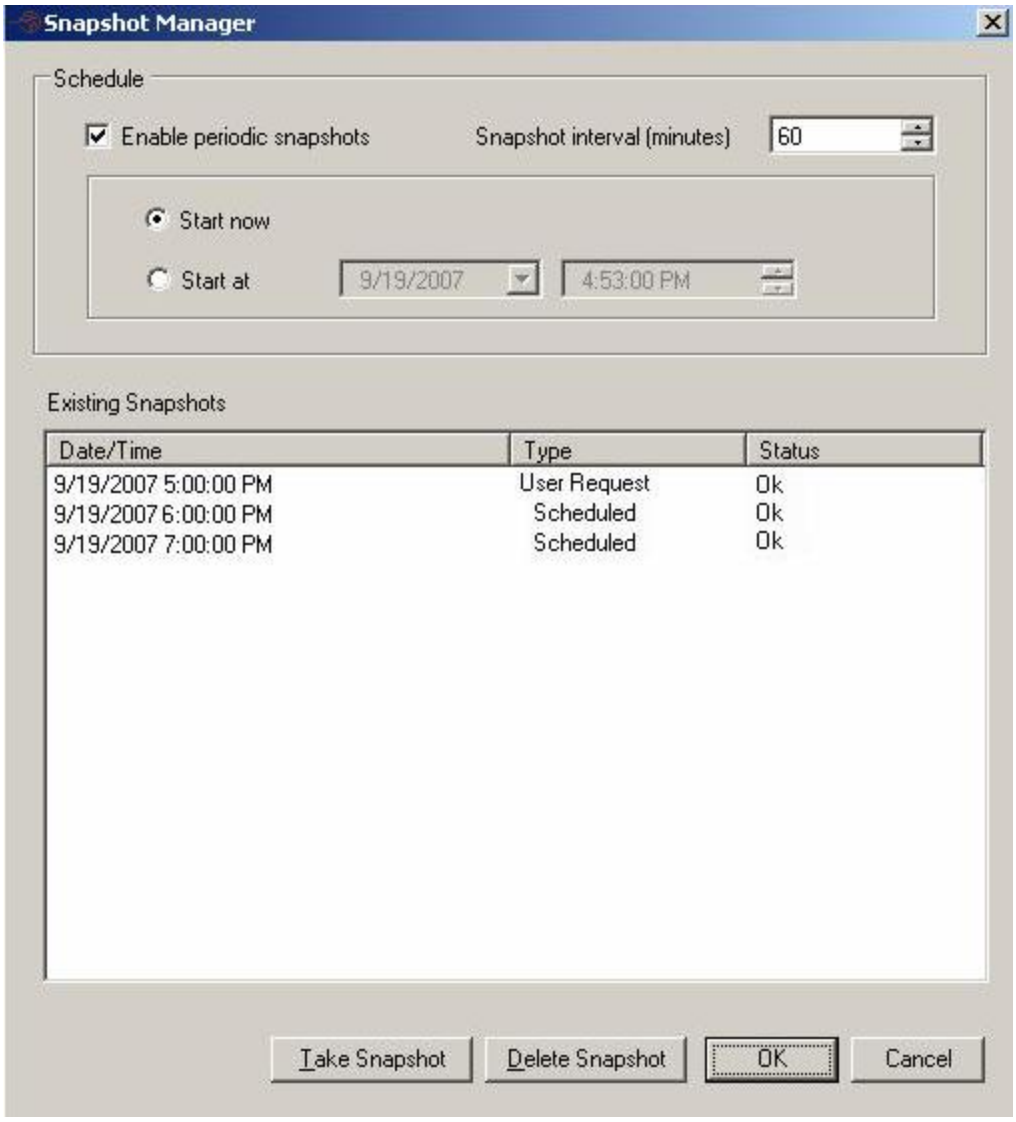

Use the following options to manage your full-server and application workload snapshots. (The dialog box will appear different between the Full-Server Failover Manager and the Application Manager, but the options are the same.)

- **Enable periodic snapshots**—By default, periodic snapshots are enabled. This option will not be available if you do not meet snapshot requirements. If you disable or do not have access to snapshots, the data on the target at the time of a failure will be used.
- **Snapshot Interval**—By default, a snapshot of the target data is taken every 60 minutes. If desired, increase or decrease the interval between snapshots.
- **Start now**—If you want to start taking snapshots immediately after the protection is established, select **Start now**.
- **Start at**—If you want to start taking snapshots at a specific data and time, select **Start at** and specify the date and time parameters.
- **Take Snapshot**—If you want to take a snapshot manually (outside of the specified interval), click **Take Snapshot**.
- **Delete Snapshot**—If you no longer want to keep a snapshot, you can delete it by highlighting the snapshot in the list and clicking **Delete Snapshot**. To help you understand the snapshots, use the **Type** and **Status** columns. The **Status** indicates the state of the connection between the source and target at the time the snapshot was taken. The **Type** indicates the kind of snapshot.
    - **Scheduled**—This snapshot was taken as part of a periodic snapshot.
    - **Deferred**—This snapshot was taken as part of a periodic snapshot, although it did not occur at the specified interval because the connection between the source and target was not in a good state.
    - **User Request**—This snapshot was taken manually by a user.

Note: The **Schedule** options at the top of the are the same options from when you configured protection. If you change the options in one location, they will be changed in the other location too.

The **Existing Snapshots** list only contains snapshots from full-server and application protection workloads. Snapshots from other utilities and tools will not be listed.

Security

To ensure protection of your data, Double-Take Software products offer multi-level security using native operating system security features. Privileges are granted through membership in user groups defined on each machine. To gain access to a source or target, the user must provide a valid operating system user name and password and the specified user name must be a member of one of the Double-Take Availability security groups. Once a valid user name and password have been provided and the source or target has verified membership in one of the security groups, the user is granted appropriate access to the source or target and the corresponding features are enabled in the client. Access is granted on one of the following three levels.

- **Administrator Access**—All features are available for that machine.
- **Monitor Access**—Servers and statistics can be viewed, but functionality is not available.
- **No Access**—Servers appear in the clients, but no access to view the server details is available.

Although passwords are encrypted when they are stored, Double-Take Software security design does assume that any machine running the client application is protected from unauthorized access. If you are running the client and step away from your machine, you must protect your machine from unauthorized access.

- Security credentials
- Adding users to the security groups
- Changing the account used to run the Double-Take service
- Configuring the Double-Take service for Active Directory

Security credentials

When a client machine attempts to access a source or target machine running on Windows, it will attempt to automatically logon to the source or target using the three methods below.

- The security credentials of the user currently logged into the client machine are sent to the source or target machine. From the security credentials, the source or target machine determines if the user is a member of the security groups and if so, grants the appropriate level of access.
- The last valid set of credentials (credentials previously granting either Administrator or Monitor level access) used to access each machine is recorded in the registry of the client machine. If the logon attempt using the credentials of the user currently logged in fails, a set of credentials is retrieved from the registry and is sent to the source or target. The source or target checks the validity of the credentials and determines if the user is a member of one of the security groups and then grants the appropriate level of access.

Note: You can disable the feature that maintains the security credentials in the registry.

- Each valid set of credentials (credentials previously granting either Administrator or Monitor level access) used by the client application is recorded in a memory-resident credentials buffer maintained by the client application. If the logon attempts using the credentials of the user currently logged in or those credentials stored in the registry fails, a set of credentials is retrieved from the client application's credentials buffer and is sent to the source or target. This process is repeated until a valid set of credentials is found or the credentials buffer is exhausted.

Note: The credentials buffer is cleared each time the client application is closed.

The client tries each of these three methods until a set of credentials granting Administrator access is found. If no credentials granting Administrator access are found, the client attempts to find a set of credentials granting Monitor access. If no credentials grant Monitor access, the user must manually logon to the source or target by providing a user name, password, and domain.

Note: If a user name exists both on the local machine and on the network, Windows first attempts to login to the machine with the local user name and password and ignores the domain. If this fails, it then tries to login with the network user name, password and domain.

Adding users to the security groups

The security groups are automatically created during the installation process. The groups are assigned specific case-sensitive names.

- **Double-Take Admin**
- **Double-Take Monitors**

The local administrator and the domain administrator are automatically added to the **Double-Take Admin** group.

Note: If Double-Take Availability is installed on a member servers, it will use the local groups. If an Active Directory user is granted access to the Active Directory **Double-Take Admin** or **Double-Take Monitors** groups, the user or domain group must also be granted access to the local Double-Take Availability groups. If Double-Take Availability is installed on a domain controller, the Active Directory group will provide sufficient access. The groups are created in the Users OU and need to stay here. If the groups are not there, users will be unable to log into Double-Take Availability on the domain controller.

Users that need administrator access to Double-Take Availability must be added to the **Double-Take Admin** group. All users that need monitor only access must be added to the **Double-Take Monitors** group. In both cases, local users, domain users, or global groups may be added to the local groups.

To add, delete, or modify users for a group, follow these steps.

1. Select **Start**, **Programs**, **Administrative Tools**, and **User Manager**. (If you are on a domain controller, select **User Manager for Domains**.)
2. Double-click the group to be modified or highlight it and select **User**, **Properties**.
3. To add local users, domain users, and/or global groups to the group, click **Add**.
4. Select the local user, domain user, and/or global group to be included in the security group.
5. Click **OK** to return to the Local Group Properties dialog box.
6. Click **OK** to return to the User Manager.
7. Exit the User Manager.

Changing the account used to run the Double-Take service

By default, the Double-Take service is configured to log on as the system account. If you want to select a specific account to run the service, use these instructions.

Note: If you are protecting full-server workloads, you cannot modify the account used to run the Double-Take service. Otherwise, the full-server protection will not function correctly.

---

1. Modify the user account that the Double-Take service is using.
    a. Open the Double-Take service settings by selecting **Start**, **Programs**, **Administrative Tools**, **Services** and double-clicking the Double-Take service.
    b. Select the **Log On** tab, select **This Account**, and enter a valid domain account.
    c. Enter the password for this account.
    d. Click **OK** to save these settings.
2. Grant an additional user right to the account you are using to run the Double-Take service.

    If domain-level policy settings are defined (through **Domain Security Policy**, **Security Settings**, **Local Policies**, **User Rights Assignment**), they will override local policy settings.

    ---

    a. Select **Start**, **Programs**, **Administrative Tools**, **Local Security Policy**.
    b. Expand the **Local Policies** folder and highlight the **User Rights Assignment** folder.
    c. Double-click the option **Act as part of operating system** on the right pane of the screen.
    d. Add the user that you selected to run the Double-Take service and click **OK**.
    e. Exit the Local Security Settings dialog box. This user is now configured to run the Double-Take service.

3. Add the domain account to the local administrator group.
    a. Select **Start**, **Programs**, **Administrative Tools**, **Computer Management**.
    b. Expand the **Local Users and Groups** folder and highlight the **Groups** folder.
    c. Right-click on the **Administrators** group on the right pane of the screen and select **Add to Group**.
    d. Click **Add**.
    e. Locate the domain account that you are using for the Double-Take service. Select that account and click **OK**.
    f. Click **OK** to close the Administrators Properties dialog box.
    g. The domain account is now added to the local administrator group. Close the Computer Management window.

Configuring the Double-Take service for Active Directory

If you want to use Active Directory registration, the Double-Take service must have privileges to modify Active Directory. There are two options for assigning the privileges.

- **User account**—Assign a user account to the Double-Take service and assign Active Directory privileges to that user. Refer to your Windows reference guide for Active Directory privileges.
- **Active Directory object**—Give the computer (or domain computers for all computers within a domain) read/write access to the **Double-Take Instances** object in Active Directory.

Use the following instructions to configure the Double-Take service.

1. Select **Start**, **Programs**, **Administrative Tools**, **Active Directory Users and Computers**.
2. Verify that **Advanced Features** is enabled on the **View** menu so that the **System** folder is displayed.
3. Expand the **System** folder and select **WinsockServices**.
    - If you have not run the Double-Take service under the domain administrator account or an account with update privileges for Active Directory, there will be no Double-Take Availability Active Directory instance listed. You will need to right-click on the **Winsock Services** folder to modify the setup for all Active Directory instances.
    - If you have run the Double-Takeservice under the domain administrator account or an account with update privileges for Active Directory, **Double-Take Instances** will be listed. You can right-click **Double-Take Instances** to modify the Active Directory setup for the one instance or right-click on the **Winsock Services** folder to modify the setup for all Active Directory instances.
4. Select the **Security** tab.
5. Click **Add** and select the specific computer account you are running Double-Take Availability on or **Domain Computers** to allow all computers within the domain to update Active Directory.
6. Grant both **Read** and **Write** access and click **OK**.

Note: If your corporate policies require that only the minimum required privileges be supplied, you can select only the permissions listed below by modifying the **Advanced** permissions for the account.

- List Contents
- Read All Properties
- Write All Properties
- Read Permissions

Evaluations

You may want to evaluate Double-Take Availability before implementing it in your production environment. This is a good process for users who want to see, first-hand, the benefits that Double-Take Availabilityhas to offer.

The evaluation processes walk you through a step-by-step process to assess the key Double-Take Availability features. Select a link below based on the type of evaluation you would like to perform.

- Evaluating data protection
- Evaluating full-server protection

For the evaluation, you should install in a test environment. Do no use actual production servers because you will be forcing a failure during the evaluation.

Evaluating data protection

The following evaluation procedure has eleven tasks containing step-by-step instructions for evaluating the data protection functionality of Double-Take Availability.

Before starting this evaluation procedure, make sure you have reviewed the server requirements and that you have installed the software on both the source and target.

Also, you should have approximately 500 MB to 1 GB of data on the source for testing. If you are going to be protecting application data, make sure the application is pre-installed on the target, but the application is not running on the target. If the application is running on the target, the files will be held open and Double-Take Availability will not be able to write to the files. In the event of a source failure, the application can be started on the target and the files can then be accessed.

Note: For the evaluation, you should install in a test environment. Do no use actual production servers because you will be forcing a failure during the evaluation.

This evaluation consists of the following tasks.

1. Establishing a connection
2. Monitoring the activity and completion of the initial mirror
3. Changing data to cause replication
4. Verifying the data changes on the target
5. Testing your target data
6. Configuring failover monitoring
7. Monitoring failover
8. Simulating a failure
9. Simulating data changes after failover
10. Initiating failback
11. Restoring your data

Establishing a connection

1. Open the Replication Console.
2. Click **Make a connection** from the right pane of the Replication Console. If that quick launch screen is no longer visible, select **Tools**, **Connection Wizard**.

Note: If the Double-Take Servers root is highlighted in the left pane of the Replication Console, the **Connection Wizard** menu option will not be available. To access the menu, expand the server tree in the left pane, and highlight a server in the tree.

---

3. The Connection Wizard opens to the Welcome screen. Review this screen and click **Next** to continue.

Note: At any time while using the Connection Wizard, click **Back** to return to previous screens and review your selections.

---

4. If you highlighted a source in the Replication Console, the source will already be selected. If it is not, select the Double-Take Availability source. This is the server that you want to protect. Click **Next** to continue.

Note: Double-Take Availability will automatically attempt to log on to the selected source using the identification of the user logged on to the local machine. If the logon is not successful, the Logon dialog box will appear prompting for your security identification. When logging in, the user name, password, and domain are limited to 100 characters.

---

5. If you highlighted a target in the Replication Console, the target will already be selected. If it is not, select the Double-Take Availability target. This is your backup server that will protect the source. Click **Next** to continue.

Double-Take Availability will automatically attempt to log on to the selected target using the identification of the user logged on to the local machine. If the logon is

not successful, the Logon dialog box will appear prompting for your security identification. When logging in, the user name, password, and domain are limited to 100 characters.

---

6. On the next screen, verify **Create a new replication set with this name** is selected.
7. Enter a name for your replication set, and click **Next** to continue.
8. A tree display appears identifying the volumes and directories available on your source. Mark the check box of the volumes and/or directories you want to protect. Click **Next** to continue.
9. There are two pre-defined locations to store the source data on the target, or you can select a custom location. For this evaluation, select the option **Send all data to the same path on the target**. This option keeps the directory structure on the source and target identical.
10. Click **Next** to continue.
11. Review your selections on the summary screen. You do not need to set any advanced options for this evaluation, so click **Finish**. The Connection Wizard will close, the connection will be established, and mirroring and replication will begin.
12. You will be prompted to save your newly created replication set. Click **Yes** to save it.

Monitoring the activity and completion of the initial mirror

View your new connection in the Replication Console by highlighting the source on the left pane. The connection will appear on the right pane. Use the horizontal scroll bar at the bottom of the right pane to view the status columns. Pay attention to the **Mirror Status** column which shows the status of the mirroring operation. During the mirroring process, you will see a percentage of the mirror that has been completed. When the **Mirror Status** changes to **Idle**, there is no mirroring activity, meaning your initial mirror has completed.

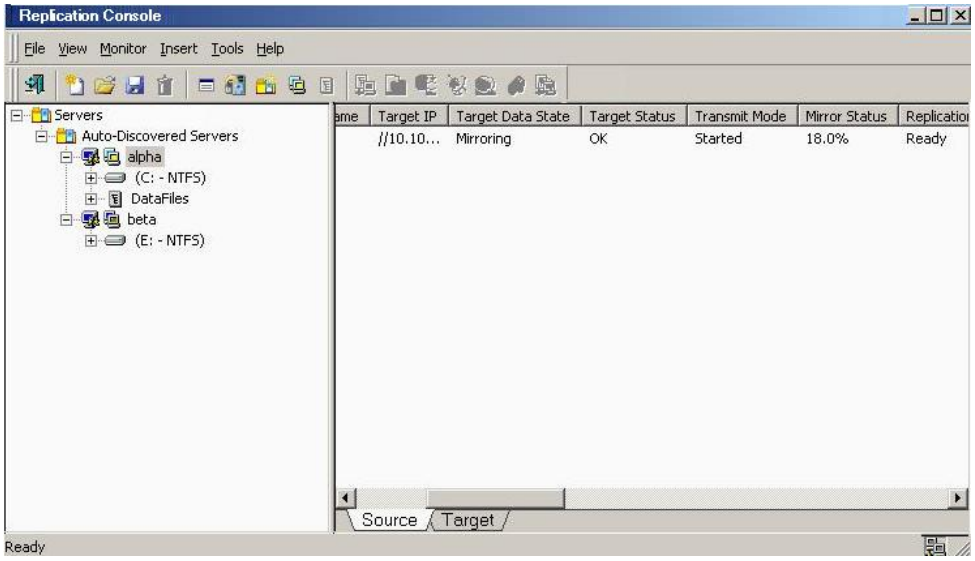

To view specific mirroring statistics that may be of interest, use the horizontal scroll bar at the bottom of the right pane to view the various columns.

- **Sent (Bytes)**—The total number of mirror and replication bytes that have been sent during this connection.
- **Sent Mirror (Bytes)**—The total number of mirror bytes only that have been sent during this connection.
- **Skipped Mirror (Bytes)**—The total number of bytes that have been skipped when performing a difference or checksum mirror. These bytes are skipped because the data is the same on the source and target machines.
- **Remaining Mirror (Bytes)**—The total number of mirror bytes only that remain to be sent to the target.

Monitoring a data workload contains complete details on all of the Replication Console statistics. After your mirror is complete, look at your target and you will see the replicated data stored in the location you specified. Now you are ready to continue with the evaluation.

Changing data to cause replication

In order to test replication, you need to change the data on your source. This includes modifying existing files, creating new files, deleting files, and changing permissions and attributes.

1. On the source, browse through the directories and files contained in your replication set.
2. Select four files from your source and record the file name, date, time, and file size for each file.
3. On your target, locate those same four files that you just identified on your source. The files on the target match the files on the source.
4. Back on your source, view the contents of one of your files contained in your replication set and note the file contents.
5. On your target, view that same file that you just viewed on the source. The file contents on the target match the file contents on the source.
6. Highlight your source in the left pane of the Replication Console.
7. Locate the **Replication Status** and **Sent (Bytes)** columns in the right pane.
8. Tile your Replication Console so that you can see it while still having access to your desktop.
9. On your source, edit the file that you viewed above. Save your changes, and watch the Replication Console statistics as the file change causes replication to occur.
10. Modify the other three files so that the date, time, and/or size is updated, and again watch the Replication Console statistics as the file changes cause replication to occur. While Double-Take Availability is actively replicating, the status will be **Replicating**. When there is no replication activity, the status is **Ready**.
11. Use the horizontal scroll bars to display additional replication statistics.
    - **Sent (Bytes)**—The total number of mirror and replication bytes that have been sent during this connection
    - **Queued Replication (Bytes)**—The total number of replication bytes that remain in the source queue
    - **Sent Replication (Bytes)**—The total number of replication bytes that have been sent during this connection
    - **Last File Touched**—Identifies the last file that Double-Take Availability transmitted to the target

Monitoring a data workload contains complete details on all of the Replication Console statistics.

Note: Many user applications typically save an entire file even though only a portion of the file may have changed. Therefore, the replication statistics will show the entire file being transmitted, not just the changed data. To confirm that replication only transmits the changed segments of files, you must use an application, such as a database application, or a command, such as the echo command, to save only the changed portions of a file.

You may notice your **Replication Status** toggle between **Replicating** and **Ready** as it continues processing the file changes, when your **Replication Status** stays at **Ready**, Double-Take Availability is waiting for additional changes to transmit. After replication is complete, you are ready to continue with the evaluation.

Verifying the data changes on the target

Now that you have modified some of the files, you want to be sure that the file modifications were applied correctly.

Note: Because of the way the Windows Cache Manager handles memory, machines that are doing minimal or light processing, as you are in this evaluation, may have file operations that remain in the cache until additional operations flush them out. This may make Double-Take Availability files on the target appear as if they are not synchronized. When the Windows Cache Manager releases the operations in the cache on the source, the files will be updated on the target. To make sure this does not impact your testing, flush the cache by copying a couple of files from one directory to another and then deleting them.

---

1. Browse your source and target. Compare the directory structures and the total number of files.
2. Look again at the four files you modified earlier. Verify manually that the changes you made have been applied to the target copy of the file.
3. Right-click the connection on the right pane of the Replication Console and select **Verify**. You will see two choices on the Start Verification dialog box.
   - **Verify only**—This option performs the verification process by comparing the date, time and size of each file and generates a verification report identifying the files that are not synchronized.
   - **Remirror data to the target automatically**—This option performs the verification process by the comparison method specified, generates a verification report, and then remirrors those files from the source to the target that are not synchronized.
4. Select **Verify only** and click **OK**.

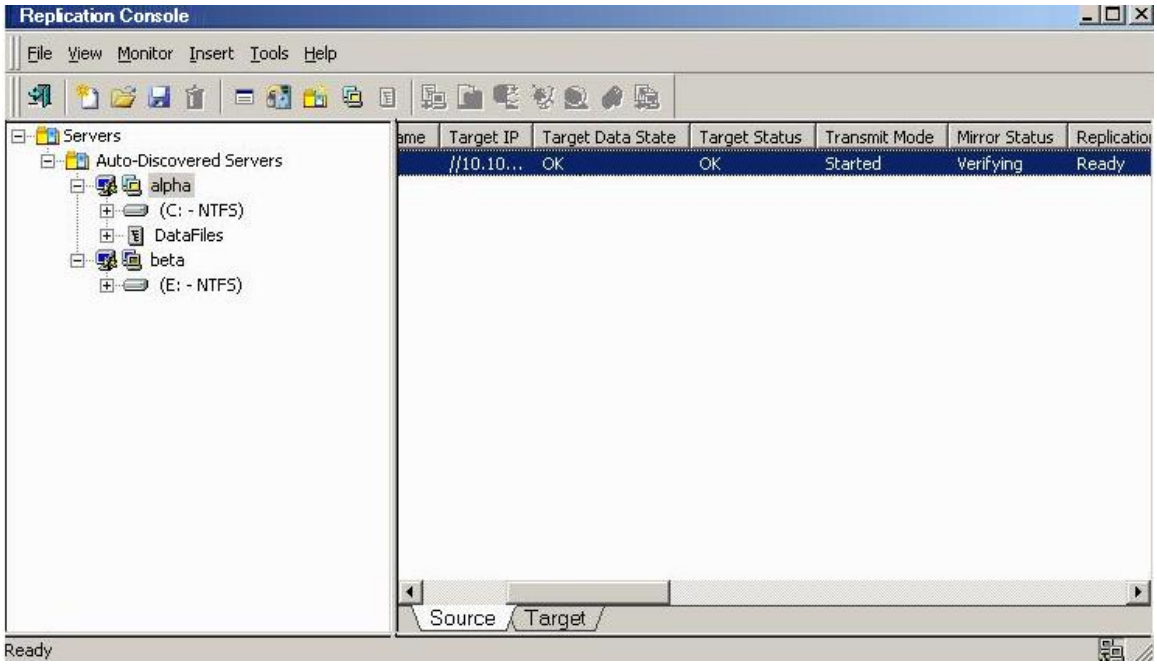

Just like when you were monitoring the mirror and replication processes, you can monitor the verification process. Notice that **the Mirror Status** column changes to **Verifying** while the verification process occurs. When the verification is complete, Double-Take Availability will have created a log file for you to review.

5. Wait until your **Mirror Status** has returned to **Idle** and then open the file DTVerify.log located in the Double-Take Availability installation directory on your source. You will see that all of the files are reported as the same.
6. Modify one of your files on the target and repeat the verification process, but this time, select **Remirror files to the Target automatically**.

Note: Since your target file is newer, make sure that **Only if Source file is newer than Target copy** is not selected.

---

7. Look at the file on the target that you modified and confirm that your changes are gone. The source version has overwritten the file on the target.

Testing your target data

At this point in your evaluation, you may want to test your target data. The type of testing you will need to perform will depend on the type of data you are protecting.

- **User data**—If you are protecting user files, you can use the associated application to open the files on the target. Open one or more of the files to test the integrity of the data. Do not save the file after you have opened it, because that will update the copy of the data on the target, which you do not want to do at this point in the evaluation.
- **Application data**—If you are protecting application data, for example a database application, you will need to use that application to test the integrity of the data and the files. Use the following instructions to test application data on the target.
    1. To test the application data on the target, you will need to start the application on the target. But Double-Take Availability requires applications to be in a standby mode in order to update files on the target. In order to meet both of those requirements, you will need to pause the target. When you pause the target, the source begins to queue the data changes that are occurring, giving you an opportunity to start the services on the target, test the data, stop the services, and then resume the target. Make sure your mirror is **Idle** and then pause the target by right-clicking the connection in the Replication Console and selecting **Pause Target**.
    2. Once the target is paused, you can start the application services on the target. Test the application data by using clients to connect to the application. For this evaluation, the clients will need to be configured to access the application from the target. In a real-world scenario, if failover has occurred, the target would be standing in for the source and the clients would still be accessing the application from the source identity.
    3. After you have completed your testing, stop the application services on the target.
    4. After the application services on the target have been stopped, you can resume your target through the Replication Console by right-clicking the connection and selecting **Resume Target**.
    5. While you were testing the application on the target, the application files were updated on the target, thus your source and target are no longer synchronized. You will need to perform a manual remirror to resynchronize the files on the source and target. Right-click the connection and select **Mirroring**, **Start**. Select a **File differences** mirror. Make sure that **Send data only if Source is newer than Target** is not selected. Since your target files are actually newer than the source (because of the testing you performed),

you do want the newer files on the target to be overwritten by the files from the source. Click **OK** to begin the mirror.

When the mirror is complete, your source and target will again be synchronized and you can continue with your evaluation.

Configuring failover monitoring

The following instructions will configure failover monitoring.

1. Open the Failover Control Center.
2. Select your target from the **Target Machine** list box.
3. Click **Login** to login to the selected target.
4. Click **Add Monitor**. The Insert Source Machine dialog box appears in front of the Monitor Settings dialog box.
5. Type in your source machine name and click **OK**. The Insert Source Machine dialog box will close and the Monitor Settings dialog box will be available for updating. This is where you configure failover monitoring.
6. Select the source to be monitored by marking the check box to the left of the source server name in the **Names to Monitor** tree.

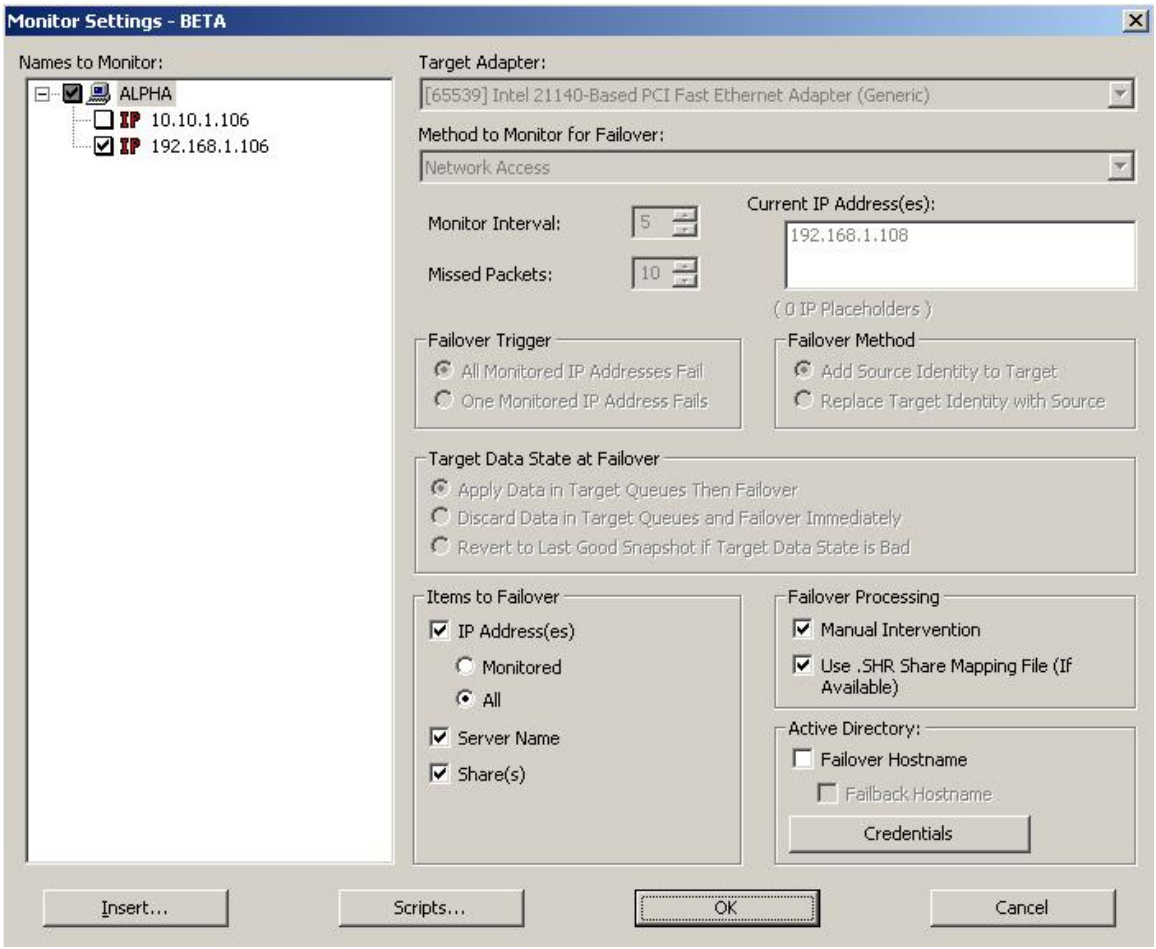

7. By default, **Failover Hostname** is disabled. This option automatically removes the host SPN (Service Principle Name) from Active Directory on the source and adds it

to Active Directory on the target. If you are using Active Directory, enable this option or you may experience problems with failover.

8. **Failback Hostname** returns the host SPN on the source and target back to their original settings on failback. If you are using Active Directory, enable this option or you may experience problems with failback.
9. If you are failing over or failing back hostnames, you need to specify an Active Directory user that has update privileges within Active Directory. Click **Credentials** and identify a user and the associated password that has privileges to create and delete SPNs. The username must be in the format fully_qualified_domain\user. The Active Directory account password cannot be blank. Click **OK** to return to the Monitor Settings dialog box.

At this point, in terms of your evaluation, your failover configuration is complete because you will be using the default settings for the remaining options. But while you are viewing the Monitor Settings dialog box, notice the flexible configuration options available to you.

Monitoring failover

Since it can be essential to quickly know the status of failover, Double-Take Availability offers various methods for monitoring the state of failover. When the Failover Control Center is running, you will see four visual indicators.

- The Failover Control Center Time to Fail counter
- The Failover Control Center status bar located at the bottom of the window
- The Failover Control Center colored bullets to the left of each IP address and source machine
- The Windows desktop icon tray containing a failover icon

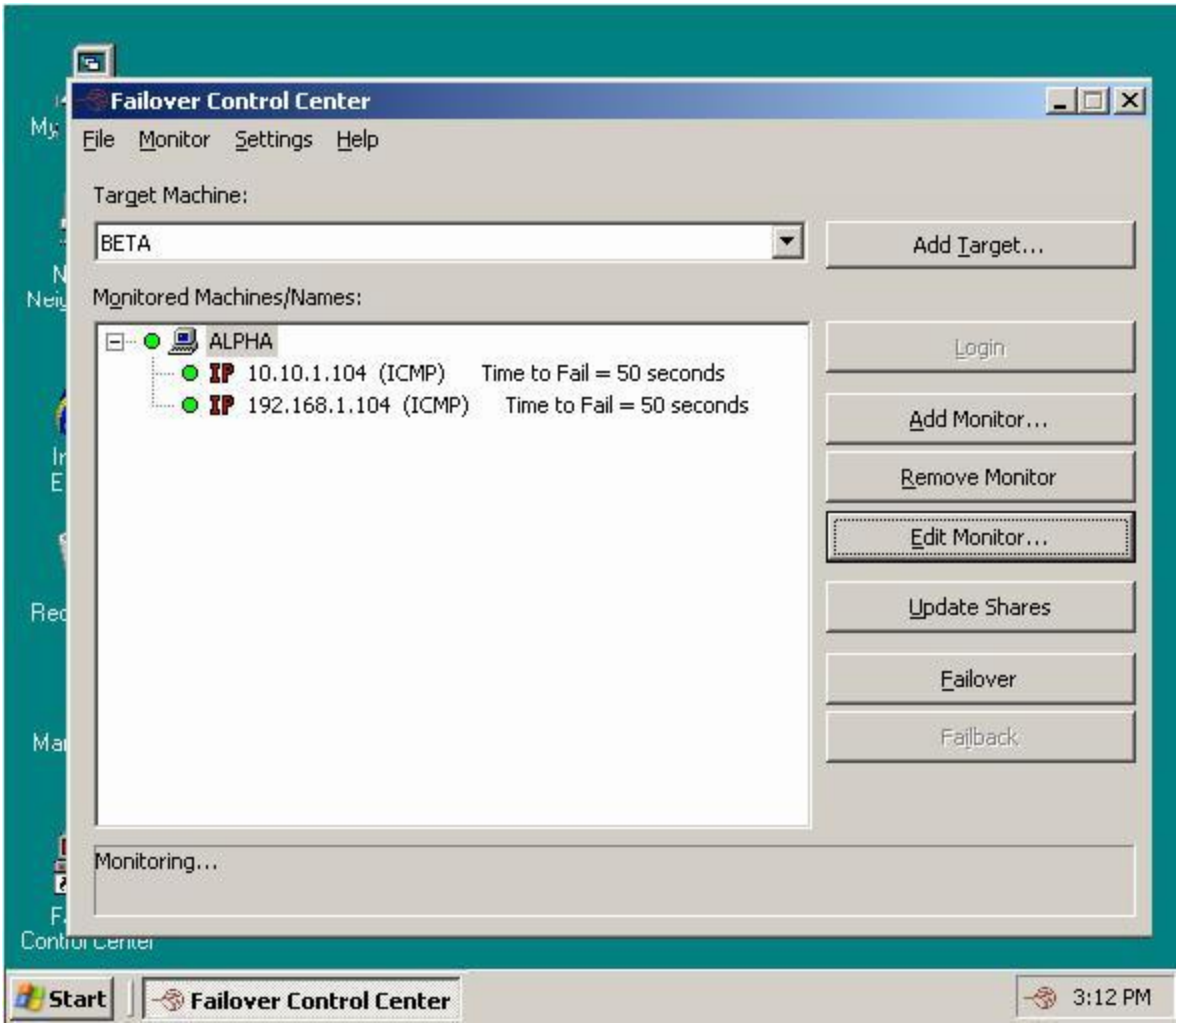

Note: You can minimize the Failover Control Center and, although it will not appear in your Windows taskbar, it will still be active and the failover icon will still appear in the desktop icon tray.

The Failover Control Center does not have to be running for failover to occur.

Monitoring failover monitoring contains more information on the Failover Control Center visual indicators.

Simulating a failure

To fully evaluate failover, you need to simulate a failure. The Failover Control Center does not have to be running in order for failover to occur, but for the purpose of this evaluation, make sure that it is running so that you can see each step of the process.

1. Ping the source's IP address from a client machine.
2. Ping the source's machine name from a client machine.
3. Disconnect the network cable(s) on the source. Notice immediately, that the Failover Control Center **Time to Fail** counter decreases and never resets. You will see the icons change to yellow and eventually to red. Once the icons are red and the **Failed Over** message is displayed, failover has occurred.
4. You will be prompted to determine how to apply the data in queue on the target. Select **Apply Data in Target Queues Then Failover**. Once the icons are red and the **Failed Over** message is displayed, failover has occurred.

Note: The Event log on the target provides details on the actual steps that have occurred during failover.

---

5. Ping the source's IP address from a client machine.
6. Ping the source's machine name from a client machine.

As you can see, the target has taken on the identity of the source. Application and user requests destined for the source are routed directly to the target. The impact on your end users is minimal.

Simulating data changes after failover

While your source is failed over to your target, end users continue to work without interruption and the data on the target will be updated. To fully evaluate the next step, restoration, simulate the changes that the end users would have made on the target while the source was unavailable.

1. Identify the file that you edited earlier on the source.
2. Locate that same file on the target and make edits to it. Save the changes.
3. Repeat that process, modifying the other three files from earlier, but this time make the modifications on the target copy of the file. Save the changes.

If desired, you can also test the target data as you did earlier. You can test user data using the associated application, and you can save the changes if desired. If you want to test application data, start the application services on the target, and test the application data by using clients to connect to the application. Because the source is now failed over, you will not need to worry about pausing the target or configuring clients to access the application from the target. The clients will continue to access the source, which is now being handled by the target machine.

Initiating failback

When failover occurs, a source machine has failed. The steps below must be used to complete failback, which releases the source identity from the target.

1. If this were a real failure scenario and not an evaluation, you would first verify that your source machine is not connected to the network. If it is, you would have to disconnect it from the network.
2. Next you would resolve the source machine problem that caused the failure.

Note: Do not connect the source machine to the network at this time.

---

3. In the Failover Control Center, select the target that is currently standing in for the failed source.
4. Select the failed source and click **Failback**.

5. You will be prompted to determine if you want to continue monitoring the source. Do not make any selections at this time.

6. At this time, you would connect the source to the network. For this evaluation, reconnect the network cable(s) on the source that you disconnected to simulate the failure.
7. After the source is online, select **Stop** in the Failover Control Center to indicate that you do not want to continue monitoring the source.

At this time, your target is back to its original identity and the source is back online.

Restoring your data

The Replication Console provides an easy method for restoring replicated data from the target back to the original source or to a new source server. You are only required to input the original source, the target, and the name of the replication set you want to restore. Double-Take Availability handles the rest, including selecting the files in the replication set and restoring them to the correct location.

1. From the Replication Console, select **Tools**, **Restoration Manager**.

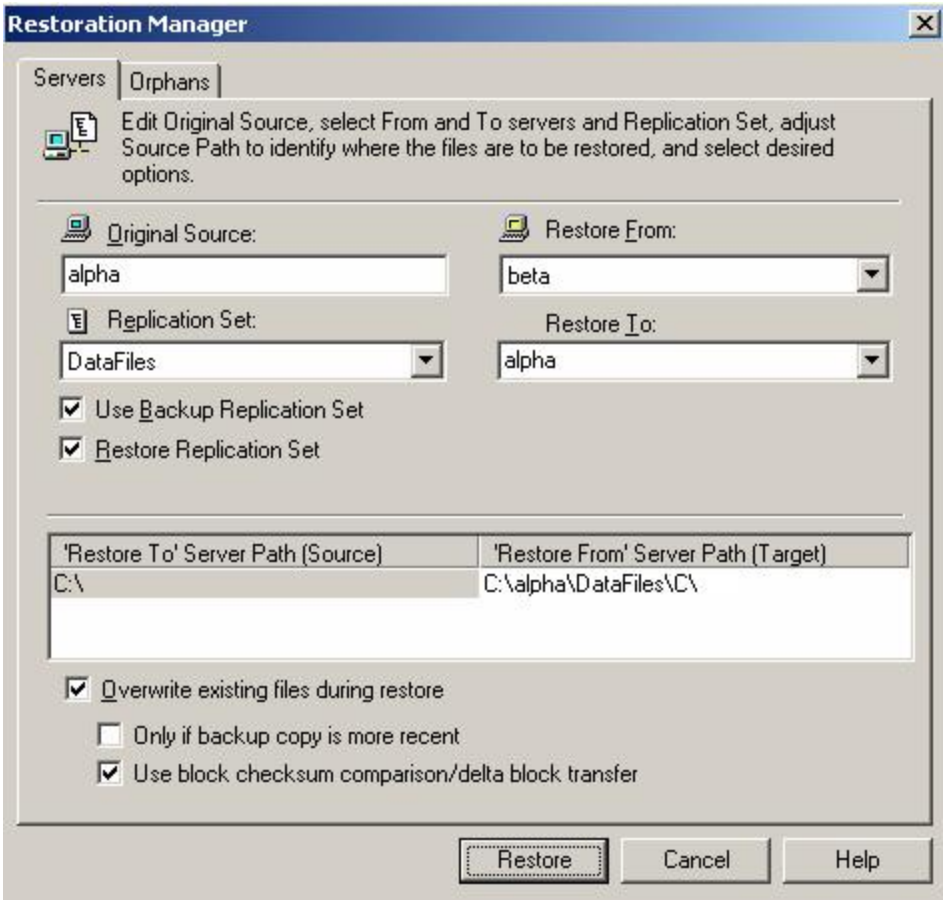

2. Identify the **Original Source** machine. This is your source machine where the data originally resided.
3. Select the **Restore From** machine. This is the target machine where the copy of the data is stored.
4. **Replication Set** contains the replication set information stored on the target machine (the machine in **Restore From**). If no replication sets are available, the list will be blank. Select the replication set that corresponds to the data that you need to restore.

5. Select the **Restore To** machine. For this evaluation, select the original source. This is the machine where the copy of the data will be sent.
6. The **Restore To** and **Restore From** paths will automatically be populated when the replication set is selected. The restore to path is the directory that is the common parent directory for all of the directories in the replication set. If the replication set crosses volumes, then there will be a separate path for each volume. The restore from path is the path on the target server where the replicated files are located.
7. Use the default settings for the remaining restoration options.
8. Click **Restore** to begin the restoration.
9. After the restoration is complete, disconnect the restoration connection. You can identify a restoration connection because it is enclosed in parenthesis ( ) and it has _Restore appended to the end of the replication set name.

Once the restoration is complete, your evaluation is complete. Congratulations!

Evaluating full-server protection

The following evaluation procedure has five tasks containing step-by-step instructions for evaluating the full-server protection functionality of Double-Take Availability.

Before starting this evaluation procedure, make sure you have reviewed the server requirements and that you have installed the software on both the source and target. Also, make sure that the target you select is compatible to stand-in as the source.

Note: For the evaluation, you should install in a test environment. Do no use actual production servers because you will be forcing a failure during the evaluation.

This evaluation consists of the following tasks.

1. Establishing full-server protection
2. Monitoring the activity and completion of the initial mirror
3. Changing data to cause replication and verifying the data changes
4. Simulating a failure
5. Starting failover

Establishing full-server protection

1. Login as a user that is a member of both the **Double-Take Admin** and local Administrators security groups.
2. Open the Full-Server Failover Manager.
3. Enter your source and target servers. You can click **Browse** when selecting either server to locate it by drilling down through your network. After you have specified a server name, enter login credentials when prompted. Once the server is selected and logged in, the **Properties** tab at the bottom of the Full-Server Failover Manager updates to display the server's properties.

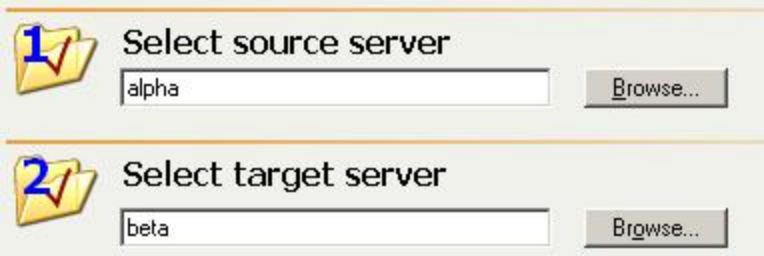

4. Optional protection settings are available but are not required for the evaluation. Feel free to review the optional settings, if desired. A complete description of each setting can be found in Optional full-server protection settings.
5. You must validate that your target is compatible with your source and can stand-in if the source fails. Click **Validate** configuration. The **Validation** tab at the bottom of the Full-Server Failover Manager updates to display the validation check. Errors are designated by a white X inside a red circle. Warnings are designated by a black exclamation point (!) inside a yellow triangle. A successful validation is designated by a white checkmark inside a green circle.

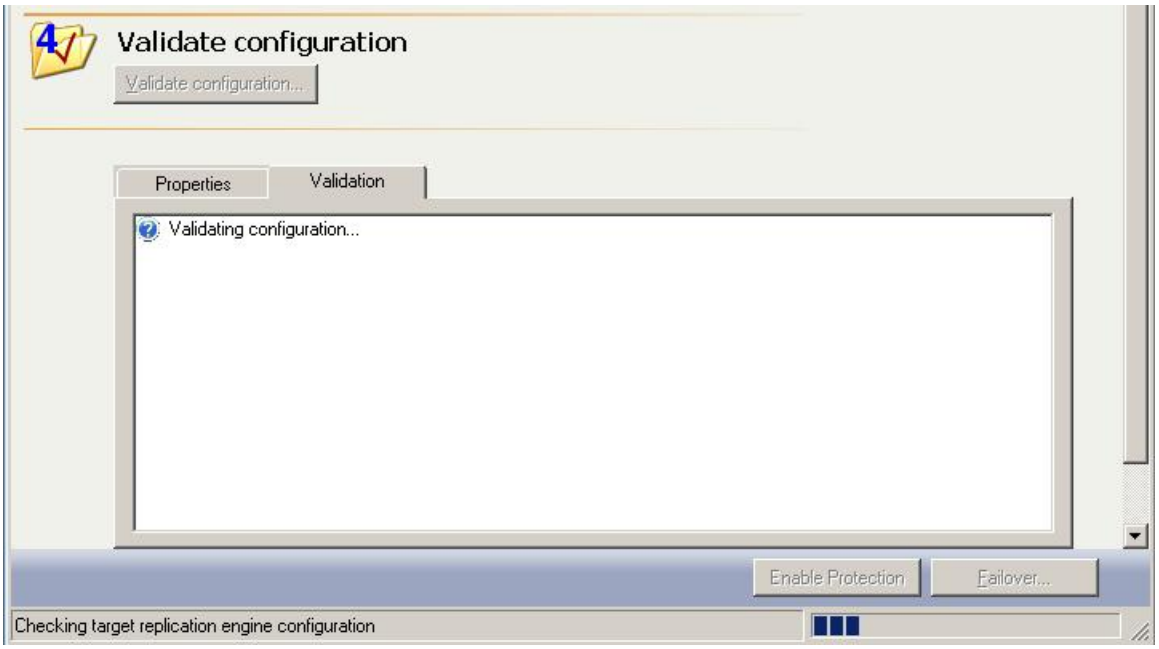

6. Double-click on any of the validation items to see details. You must correct any errors before you can enable protection. Depending on the error, you may be able to click **Fix** or **Fix All** and let Double-Take Availability correct the problem for you. For those errors that Double-Take Availability cannot correct automatically, you will need to modify the target to correct the error, or you can select a different target. You must revalidate the selected servers until the validation check passes without errors.
7. Once the validation check passes without errors, click **Enable Protection** to begin monitoring.

Monitoring the activity and completion of the initial mirror

After you have enabled full-server protection, you can monitor the protection from the Full-Server Failover Manager. The **Protection Status** is displayed in the right center of the Full-Server Failover Manager. You can tell the status of your protection from this field.

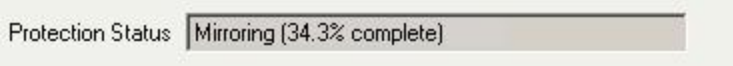

1. Watch as the mirroring percentage increases. When the mirroring is complete, the status will change to **Enabled**.
2. Look at your target and you will see the data from the source.

Once protection is **Enabled**, if the source should fail, the target can stand-in for the source.

Changing data to cause replication and verifying the data changes

In order to test replication, you need to change the data on your source. This includes modifying existing files, creating new files, deleting files, and changing permissions and attributes.

1. On the source, browse through the directories and files.
2. Select four data files from your source and record the file name, date, time, and file size for each file.
3. On your target, locate those same four files that you just identified on your source. The files on the target match the files on the source.
4. Back on your source, view the contents of one of your data files and note the file contents.
5. On your target, view that same file that you just viewed on the source. The file contents on the target match the file contents on the source.
6. On your source, edit the file that you viewed above and save your changes.
7. Modify the other three files so that the date, time, and/or size is updated.
8. Verify the changes that you made on the source have been applied to the target copy of the file.

Note: Because of the way the Windows Cache Manager handles memory, machines that are doing minimal or light processing, as you are in this evaluation, may have file operations that remain in the cache until additional operations flush them out. This may make Double-Take Availability files on the target appear as if they are not synchronized. When the Windows Cache Manager releases the operations in the cache on the source, the files will be updated on the target. To make sure this does not impact your testing, flush the cache by copying a couple of files from one directory to another and then deleting them.

Simulating a failure

To fully evaluate failover, you need to simulate a failure. To do this, power off the source. Watch as the status changes to **Failover condition met**.

Starting failover

When a failover condition is met, you will want to start failover. You can actually start failover without a failover condition, as long as protection is enabled. For example, you may want to force a failover when upgrading to a better source server.

Note: If you are testing failover and your source is a domain controller, do not let the domain controller communicate with any other production domain controllers after failover. Otherwise, when the original source domain controller is brought online after the test, it may create a USN rollback scenario if the test domain controller was allowed to communicate with other production domain controllers.

---

1. To begin failover, click **Failover**.
2. If Double-Take Availability determines there is a possibility that the data on the target is incomplete, you will be warned before failover begins. If you proceed with failover, the state of the source will be unknown until failover is complete. The best case scenario would be a missing data file, while the worst case scenario would be missing system state data that causes the server to be unusable or unbootable. For this evaluation, select **Use live data** and click **OK**.

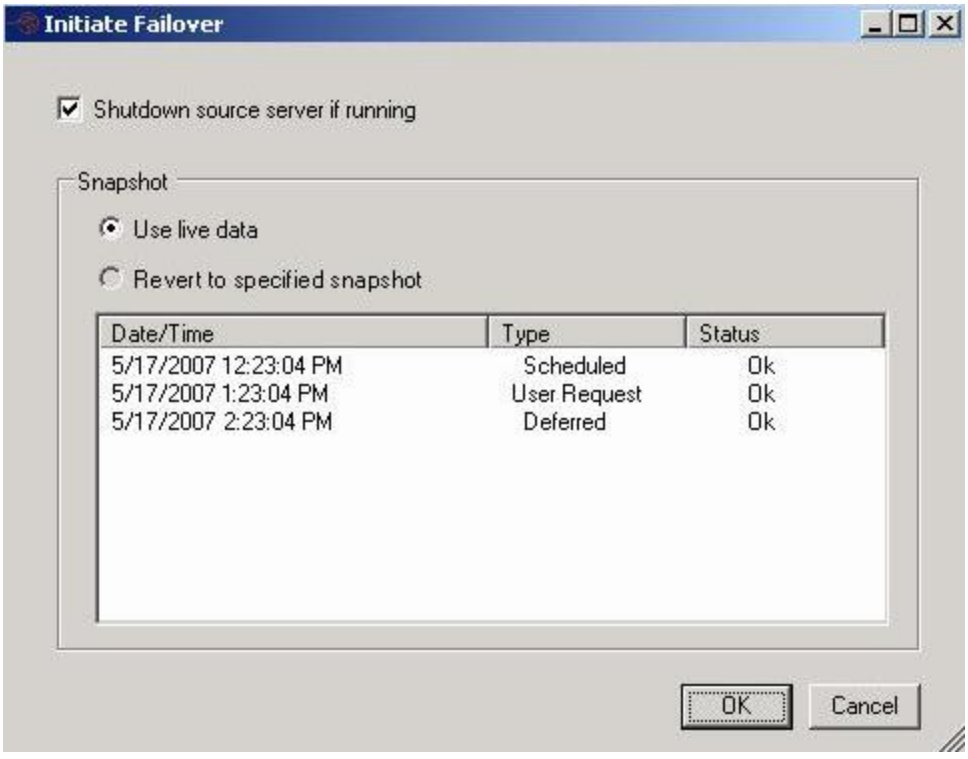

3. Monitor the failover percentage as shown in the **Protection Status**. At the end of

failover, the target will be rebooted automatically. After the reboot, the target will no longer exist, since it will become the source.

Note: Because the Windows product activation is dependent on hardware, you may need to reactivate your Windows registration after failover. Follow the on-screen prompts to complete the reactivation.

After your target has failed over and becomes your source, you can stay with that configuration long term. However, in some instances, it may be necessary or desired to go back to using the original hardware after you have failed over.

Index

B

C

D

E

F

N

O

P

Q

R

S

T

U

V

W