Double-Take®AVAILABILITY™

Version 4.7.1
User's Guide



VISION SOLUTIONS®

Contents

Chapter 1 Double-Take Availability overview **6**
Core operations 7
Mirroring 7
Replication 8
Failover 9
Restoration 10
Full-server protection 11
Supported configurations 12
One-to-one, active/standby 12
One-to-one, active/active 13
Many-to-one 14
One-to-many 15
Chained 16

Chapter 2 Double-Take Availability requirements **17**
Source and target requirements 17
Console requirements 21
Documentation requirements 21

Chapter 3 Installation **22**
Installation and upgrade notes 23
Installing or upgrading Double-Take Availability Linux servers 24
Installing or upgrading Double-Take Availability Windows client 26
Removing (erasing) Double-Take Availability 27

Chapter 4 DTSetup **28**
Running DTSetup 29
Setup tasks 30
Activating your server 30
Modifying security groups 30
Replication configuration 31
Configuring file system replication 31
Configuring block device replication 33
Configuring full-server protection 35
Configuring server settings 36
Configuring driver performance settings 36
Starting and stopping the daemon 38
Starting DTCL 39
Viewing documentation and troubleshooting tools 40
DTSetup menus 41

Chapter 5 Clients **42**
Replication Console 43
Using Replication Console workspaces 43
Clearing stored security credentials 44

Failover Control Center 45
Setting the frequency of Failover Control Center console refreshes 45

Chapter 6 Data protection 46
Establishing a data connection using the automated Connection Wizard 47
Creating a replication set 50
Establishing a connection manually using the Connection Manager 53
Establishing a connection across a NAT or firewall 58
Simulating a connection 61

Chapter 7 Establishing full-server protection 62

Chapter 8 Protection monitoring 65
Monitoring a data workload 66
Connection statistics 66
Connection and sever display 70
Log files 73
Viewing the log files through a text editor 74
Viewing the Double-Take Availability log file through the Replication Console 75
Configuring the properties of the Double-Take Availability log file 77
Double-Take Availability log messages 77
Monitoring the Linux system log 85
E-mailing system messages 98
Statistics 101
Configuring the properties of the statistics file 101
Viewing the statistics file 102
Statistics 103
SNMP 111
Configuring SNMP on your server 111
SNMP traps 112
SNMP statistics 115

Chapter 9 Connections 119
Data queues 120
Queuing data 122
Auto-disconnect and auto-reconnect 125
Reconnecting automatically 127
Pausing and resuming target processing 128
Disconnecting a connection 129

Chapter 10 Mirroring 130
Stopping, starting, pausing, or resuming mirroring 131
File difference mirror options compared 132
Mirroring automatically 133
Removing orphan files 135

Chapter 11 Replication 137
Replication capabilities 138
Replication sets 142
Creating a replication set 144

Creating or modifying replication rules manually ... 146
Selecting a block device for replication ... 148
Modifying a replication set ... 149
Renaming and copying a replication set ... 149
Calculating replication set size ... 150
Exporting and importing a replication set ... 151
Deleting a replication set ... 151
Starting replication ... 152
Inserting tasks during replication ... 153

Chapter 12 Verification ... 154
Verifying manually ... 155
Verifying on a schedule ... 157
Configuring the verification log ... 159

Chapter 13 Data transmission ... 160
Stopping, starting, pausing, and resuming transmission ... 161
Scheduling data transmission ... 161
Limiting transmission bandwidth ... 167
Compressing data for transmission ... 169

Chapter 14 Failover and failback ... 171
Configuring failover monitoring ... 172
WAN considerations ... 175
Protecting NFS shares ... 177
Protecting Samba shares ... 178
Editing failover monitoring configuration ... 179
Monitoring failover monitoring ... 180
Failing over ... 183
Removing failover monitoring configuration ... 183

Chapter 15 Failback and restoration ... 184
Restoring then failing back ... 185
Failing back then restoring ... 190

Chapter 16 Restoration ... 193
Restoring files only ... 194
Restoring an entire server ... 198
Restoring an entire server to an ESX virtual server ... 202
Configuring Red Hat-based guests after restoring ... 205
Configuring Novell-based guests after restoring ... 206

Chapter 17 Server settings ... 208
Identifying a server ... 209
Licensing a server ... 211
Configuring server startup options ... 214
Configuring network communication properties for a server ... 216
Queuing data ... 218
Configuring source data processing options ... 221
Configuring target data processing options ... 223

Specifying the Double-Take Availability database storage files ... 225
Specifying file names for logging and statistics ... 227
E-mailing system messages ... 229

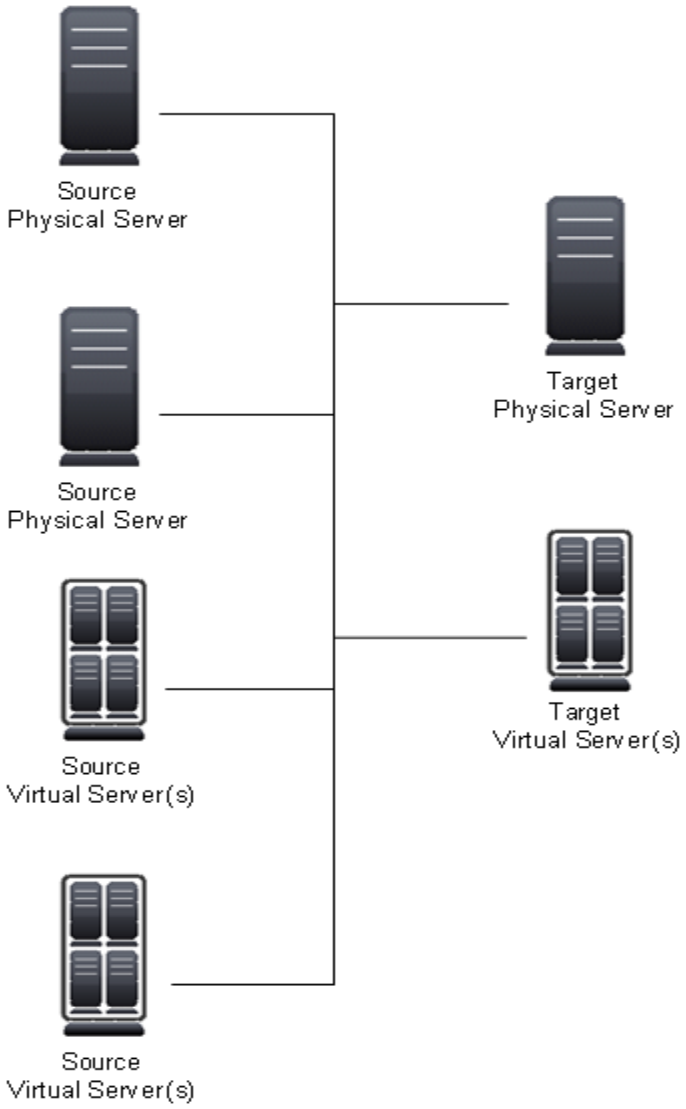
Chapter 18 Security ... 232
Logging on and off ... 233

Chapter 19 Evaluating Double-Take Availability ... 235
Establishing a connection ... 236
Monitoring the activity and completion of the initial mirror ... 238
Changing data to cause replication ... 240
Verifying the data changes on the target ... 242
Testing your target data ... 244
Configuring failover monitoring ... 245
Monitoring failover ... 246
Simulating a failure ... 248
Simulating data changes after failover ... 249
Initiating failback ... 250
Restoring your data ... 252

Chapter 1 Double-Take Availability overview

Double-Take Availability is a real-time data replication and failover software product. It augments your existing data protection strategy by reducing downtime and data loss, and it provides these services with minimal impact on existing network and communication resources.

Double-Take Availability allows you to identify and protect workloads by replicating, in real-time, data from a production server, known as the source, 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 data from the source. Double-Take Availability monitors any changes to the data and sends the changes to 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

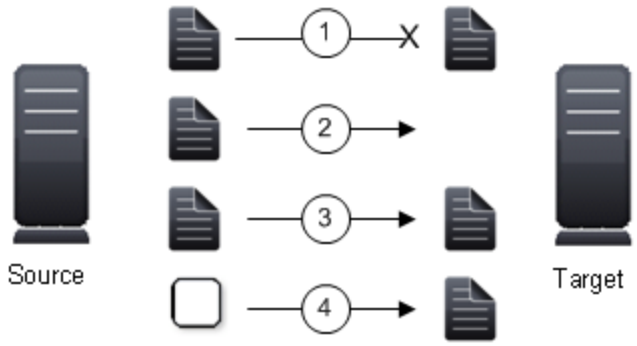
Double-Take Availability performs four basic types of operations.

- Mirroring—The initial copy or subsequent resynchronization of selected data
- Replication—The on-going capture of byte-level file changes
- Failover—The ability to 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.

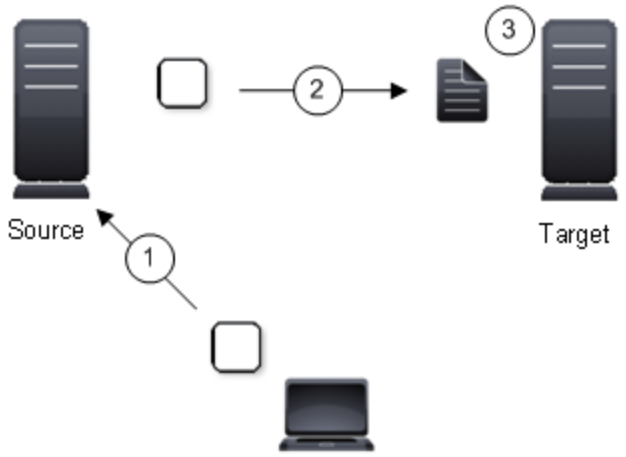


1. Identical files are not mirrored.
2. New files are mirrored.
3. Different files can be mirrored.
4. Checksums can calculate blocks of data to be mirrored.

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.



1. A user or application updates part of a file.
2. Only the changed portion of the file is replicated to the target.
3. An up-to-date copy of the file is maintained on the target.

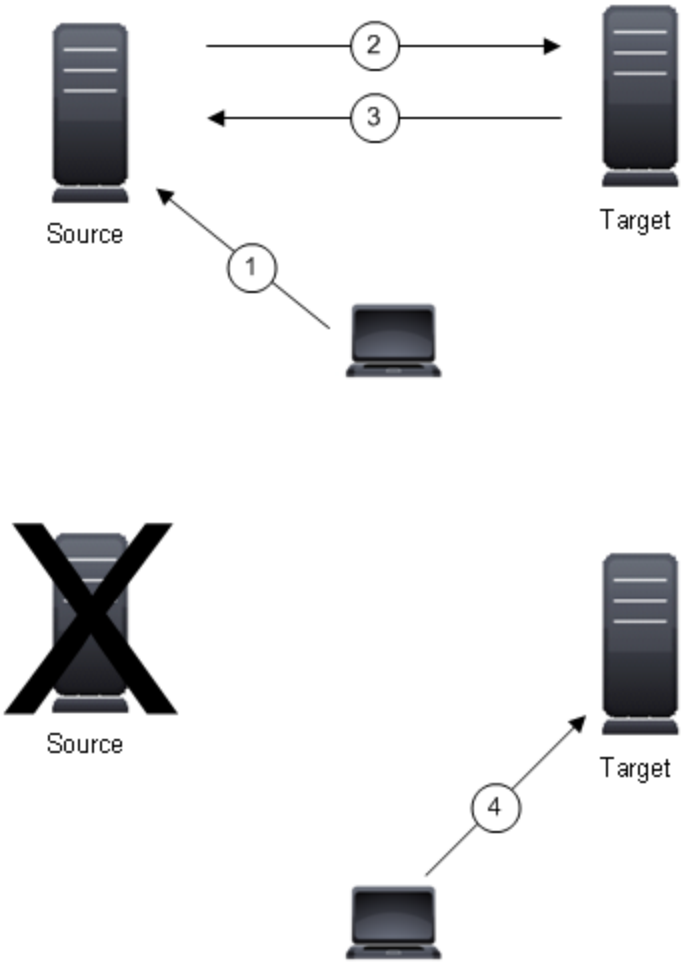
Failover

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.



1. User and application requests are sent to the source name or IP address.
2. Data on the source is mirrored and replicated to the target.
3. The target monitors the source for failure.
4. In the event the source fails, the target stands in for the source. User and application requests are still sent to the source name or IP address, which are now running on the target.

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.

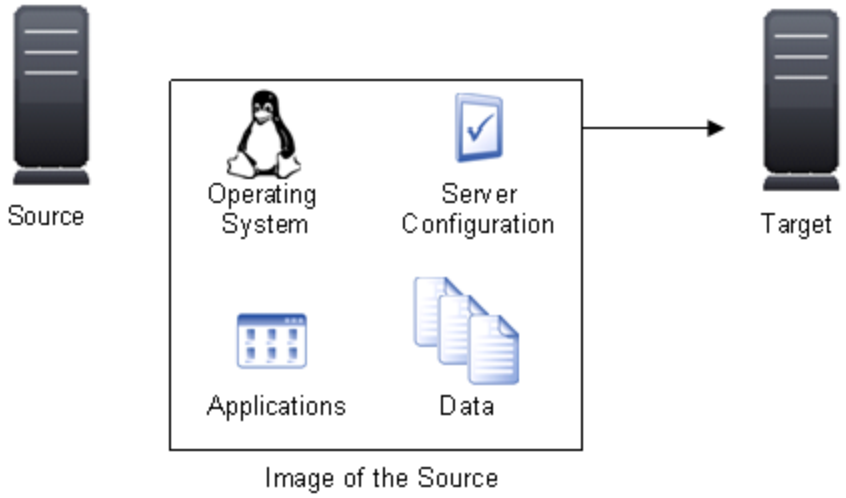
Restoration can be used if the source data is lost due to a disk crash or when the most up-to-date data exists on the target due to failover. At the time of a source server failure, your Double-Take Availability target will contain the same data as your Double-Take Availability source. If you are using 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.



1. Identical files are not restored.
2. New files are restored.
3. Different files can be restored.
4. Checksums can calculate blocks of data to be restored.

Full-server protection

In addition to protecting your data, you can also protect an entire server. An image of the entire source is replicated to the target. The image contains all of the source server's data and the source server's system state, which is the server's configured operating system and the installed applications. As changes are made on the source, Double-Take Availability mirroring and replication keeps the image of the source stored on the target up-to-date. In the event of a source failure, the image of the source stored on the target can be used to re-create the source.



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.

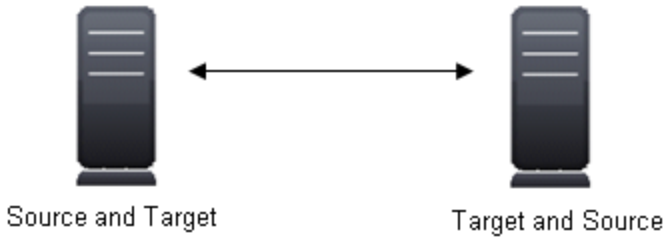
Applications

- 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

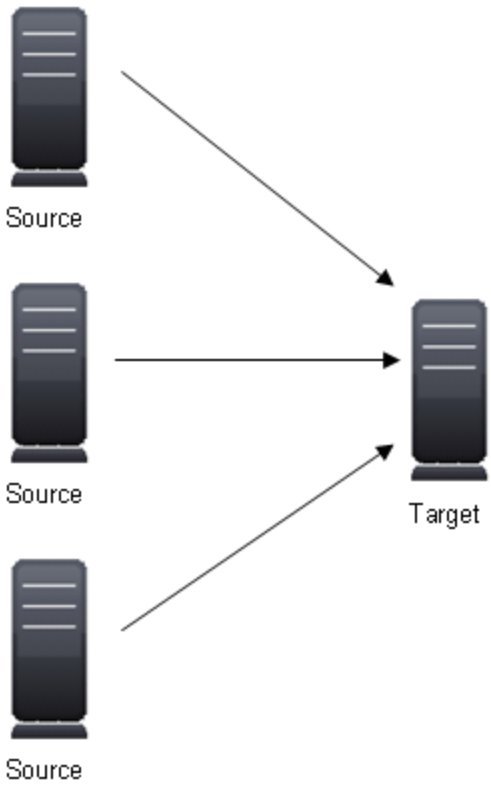
Applications

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.

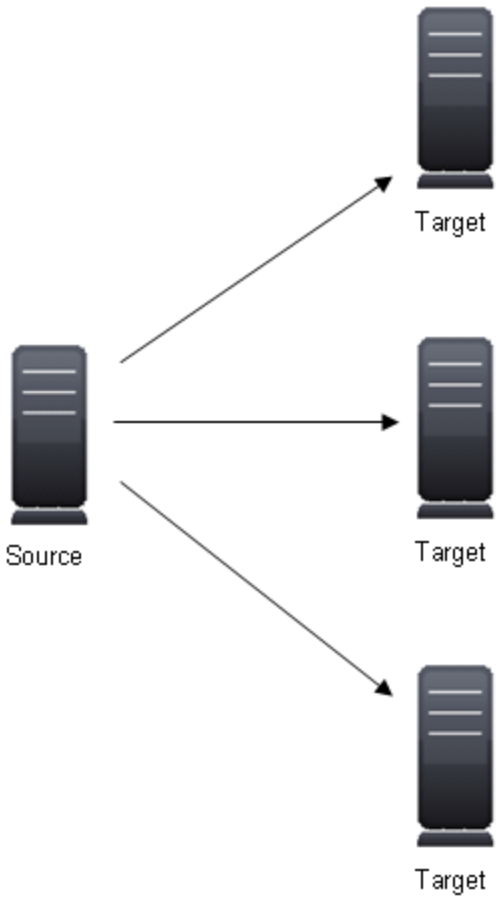
Applications

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.

Chapter 2 Double-Take Availability requirements

Each Double-Take Availability server must meet minimum requirements. Verify that each server meets the requirements for the function of that machine. Additionally, the machine where you will be running the console must also meet some basic requirements.

- Source and target server requirements
- Console requirements
- Documentation requirements

Source and target requirements

- **Operating system**—Make sure your servers meet the operating system, kernel, and file system requirements.

Red Hat Enterprise Linux 4.7 through 4.8

Kernel version—2.6.9.x.x

Kernel type for 32-bit architectures—Default, SMP, HugeMem

Kernel type for 64-bit architectures—Default, SMP, LargeSMP

File system—Ext2, Ext3, XFS

Virtual systems platforms—VMware, Microsoft Virtual Server, Xen, Virtual Iron

Red Hat Enterprise Linux 5.5 through 5.6

Kernel version—2.6.18.x.x

Kernel type for 32-bit architectures—Default (SMP), Xen, PAE

Kernel type for 64-bit architectures—Default (SMP), Xen

File system—Ext2, Ext3, Ext4, XFS

Virtual systems platforms—VMware, Microsoft Virtual Server, Xen, Virtual Iron

Red Hat Enterprise Linux 6.0

Kernel version—2.6.32.x.x

Kernel type for 32-bit architectures—Default

Kernel type for 64-bit architectures—Default

File system—Ext2, Ext3, Ext4

Virtual systems platforms—VMware, Microsoft Virtual Server, Xen, Virtual Iron, KVM

Oracle Enterprise Linux 5.5

Kernel version—2.6.18.x.x

Kernel type for 32-bit architectures—Default (SMP), Xen, PAE

Kernel type for 64-bit architectures—Default (SMP), Xen

File system—Ext2, Ext3, Ext4 (for version 5.5 only), XFS

Virtual systems platforms—VMware, Microsoft Virtual Server, Xen, Virtual Iron

CentOS 4.7 through 4.8

Kernel version—2.6.9.x.x

Kernel type for 32-bit architectures—Default, SMP, HugeMem

Kernel type for 64-bit architectures—Default, SMP, LargeSMP

File system—Ext2, Ext3, XFS

Virtual systems platforms—VMware, Microsoft Virtual Server, Xen, Virtual Iron

CentOS 5.4 through 5.5

Kernel version—2.6.18.x.x

Kernel type for 32-bit architectures—Default (SMP), Xen, PAE

Kernel type for 64-bit architectures—Default (SMP), Xen

File system—Ext2, Ext3, Ext4 (for version 5.5 only), XFS

Virtual systems platforms—VMware, Microsoft Virtual Server, Xen, Virtual Iron

SUSE Linux Enterprise Server 9.3 and 9.4

Kernel version—2.6.5-7.x

Kernel type for 32-bit architectures—Default, SMP, BigSMP

Kernel type for 64-bit architectures—Default, SMP

File system—Ext2, Ext3, ReiserFS, XFS

Virtual systems platforms—VMware, Microsoft Virtual Server

SUSE Linux Enterprise 10.2 and 10.3

Kernel version—2.6.16.xx

Kernel type for 32-bit architectures—Default, SMP, BigSMP, Xen, XenPAE

Kernel type for 64-bit architectures—Default, SMP, Xen

File system—Ext2, Ext3, ReiserFS, XFS

Virtual systems platforms—VMware, Microsoft Virtual Server, Xen, Virtual Iron

SUSE Linux Enterprise 11.0 and 11.1

Kernel version—2.6.27.19

Kernel type for 32-bit architectures—Default, Xen, XenPAE, VMI

Kernel type for 64-bit architectures—Default, Xen

File system—Ext2, Ext3, ReiserFS, XFS

Virtual systems platforms—VMware, Microsoft Virtual Server, Xen, Virtual Iron



The kernel version must match the expected kernel for the specified release version. For example, if /etc/redhat-release declares the system to be a Redhat 5.5 system, the kernel that is installed must match that.

The Standard Double-Take Availability license protects one virtual machine. The Virtual Guest 5-pack license can protect up to five virtual machines. The Premium license can protect an unlimited number of virtual machines.

- **System Memory**—At a minimum, you should have 512 MB of memory on each server. Ideally, you should have 1 GB.
- **Disk Usage**—The amount of disk space required for the Double-Take Availability program files is approximately 85 MB. About 45 MB will be located on your /(root) partition, and the remainder will be on your /usr partition. 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.

- **Protocols**—Your servers must have TCP/IP. IPv4 is the only supported version.
- **IP address and subnet configuration**—Because of limitations in the way the Linux kernel handles IP address aliases, do not mix subnets on the eth0 network interface. Failover should not cause problems in this configuration, but you will lose IP addresses during failback. Therefore, if you must mix subnets on a single interface, use eth1 or higher.
- **Name resolution**—Your servers must have name resolution or DNS. For details on name resolution options, see your Linux documentation or online Linux resources.
- **Security**—Double-Take Availability security is granted through membership in user groups. The groups can be local or LDAP (Lightweight Directory Access Protocol). A user must provide a valid local account that is a member of the Double-Take Availability security groups.
- **SELinux policy**—SELinux should be disabled or set to permissive mode on the source and target.
- **VMware Tools**—Any VMWare guest running Double-Take Availability should have the appropriate VMWare Tools package installed.

Console requirements

The Replication Console can be run on any of the following operating systems.

- Windows 2008
- Windows 2003
- Windows 7
- Windows Vista
- Windows XP Service Pack 2 or later

Documentation requirements

In order to view the online documentation, you must have one of following browsers.

- Internet Explorer 5.5 or later
- Mozilla Firefox 1.0 or later

Chapter 3 Installation

Review Double-Take Availability requirements and Installation and upgrade notes then use the appropriate instructions from the following list to meet your goal.

- Installing or upgrading Double-Take Availability Linux servers—Use these instructions if you are installing or upgrading on a Linux server.
- Completing installation and configuring replication using DTSetup—Use these instructions to configure your Double-Take Availability Linux servers.
- Installing or upgrading Double-Take Availability Windows client—Use these instructions if you want to install or upgrade on a Windows client.

Installation and upgrade notes

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

- 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.
- A script called DTInfo.sh is included in the installation. This script can be run via DTSetup to collect configuration data for use when reporting problems to technical support. Depending on the type of diagnostic information gathering you select, DTInfo can gather Double-Take Availability log files, Double-Take Availability and system settings, network configuration information, and other data which may be necessary for technical support to troubleshoot issues. After running the script, a .tar.gz is automatically created with the information gathered. You must have root (or uid 0 equivalent) to execute the diagnostics or to copy or read the resulting file.
- Double-Take Availability 4.7 is interoperable back to version 4.5 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.
  - **4.5 client**—Supports 4.5 source and target, but does not support 4.6 or 4.7 source or target
  - **4.6 client**—Supports 4.5 or 4.6 source and target as long as the target is the same or newer than the source, but does not support 4.7 source or target
  - **4.7 client**—Supports 4.5, 4.6, or 4.7 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 and the target daemon is restarted, 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.

Installing or upgrading Double-Take Availability Linux servers

Use these instructions if you are installing or upgrading Double-Take Availability on a Linux server.

1. Determine the installation package that is appropriate for your operating system. There are separate .rpm installation packages for 32-bit and 64-bit architectures. Make sure you are installing the correct .rpm file. If you are uncertain about the architecture of your machine, you can use the uname -m command to determine it. Additionally, the version_numbers in the installation file name will vary and will correspond to the version of Double-Take Availability you are installing. For example, if you are installing version 4.7.0.0.2200, the installation files would be DoubleTake-4.7.0.0-2200.i386.rpm or DoubleTake-4.7.0.0-2200.x86_64.rpm.
    - 32-bit architecture—DoubleTake-version_numbers.i386.rpm
    - 64-bit architecture—DoubleTake-version_numbers.x86_64.rpm
2. If you are upgrading an existing Double-Take Availability server, you should complete the following steps before upgrading.
    - Shutdown the protected application(s)
    - Stop the Double-Take daemon
    - Unmount DTFS or detach DTLOOP
3. Once you have determined the appropriate installation package to use, you can install the software from the UI or from the command line.
    - **UI installation**—Double-click the .rpm file from the UI and confirm the installation or upgrade.
    - **Command line installation**—Use the following steps to install from a command line.
        a. Make sure you are a root/uid 0 user.
        b. Go to a shell prompt by launching a terminal application from your UI or logging in via the Linux virtual console.
        c. If you are installing from a CD, mount the file ISO 9660 or UDF file system.
        d. To run the installation use rpm -i with the installation file name to install the software or rpm -U with the installation file name to upgrade the software. For example, if you were installing on a 32-bit operating system, you would use the command rpm -i DoubleTake-4.7.0.0-2200.i386.rpm.

           A successful installation returns you to the shell prompt. If you receive an error message during the installation, you will need to reinstall the

software. If you are unable to resolve the error, contact technical support.

Use DTSetup to configure your servers for Double-Take Availability and restart the daemon if you upgraded.

Installing or upgrading Double-Take Availability Windows client

Use the instructions to install or upgrade the Double-Take Availability client on a Microsoft Windows machine.

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>:\DblTake\i386\setup.exe.
   - **Web download**—Launch the .exe file that you downloaded from the web. When prompted, unzip the files and launch the installation.
3. When the Double-Take Availability installation program begins, the Welcome screen is displayed. Click **Next** to continue.
4. Review and accept the Vision Solutions license agreement to continue with the installation program. Click **Next** to continue.
5. Select the folder where you would like to install the Double-Take Availability clients and click **Next** to continue.
6. When you are ready to begin copying the files, click **Install**.
7. After the files have completed copying, click **Finish** to exit the installation program.

Removing (erasing) Double-Take Availability

Use these instructions if you want to remove (erase) an existing Double-Take Availability installation.

1. If you are protecting an entire server, use DTSetup to disable full-server protection. After disabling full-server protection, you will need to reboot your server before continuing with removing Double-Take Availability.
2. If you are only protecting specific data (files and folders) and not the entire server, you only need to unmount DTFS.
3. Make sure you are a root/uid 0 user.

A successful removal returns you to the shell prompt. If you receive an error message during the removal, you will need to erase the software again. If you are unable to resolve the error, contact technical support.



After completing the removal process, some files may not be completely removed until the server has been rebooted.

Chapter 4 DTSetup

DTSetup is a menu-driven application that provides easy access to Double-Take Availability server configuration. Select a link for more information on DTSetup configuration tasks.

- Running DTSetup—This topic includes instructions for launching DTSetup.
- Setup tasks—The setup tasks allow you to configure activation codes, security groups, replication configuration, server configuration, and driver performance settings.
- Daemon—Built-in scripts allow you to quickly and easily start and stop the Double-Take daemon.
- DTCL—You can launch the Double-Take Availability interactive command prompt which allows you to enter DTCL commands one at a time.
- Documentation and troubleshooting—DTSetup provides easy access to Double-Take Availability log files, a diagnostic collection tool, user documentation, and several legal documents.
- DTSetup menus—This topic includes a list overview of the DTSetup menu system. Reference the links in the list for complete details on completing tasks in DTSetup.

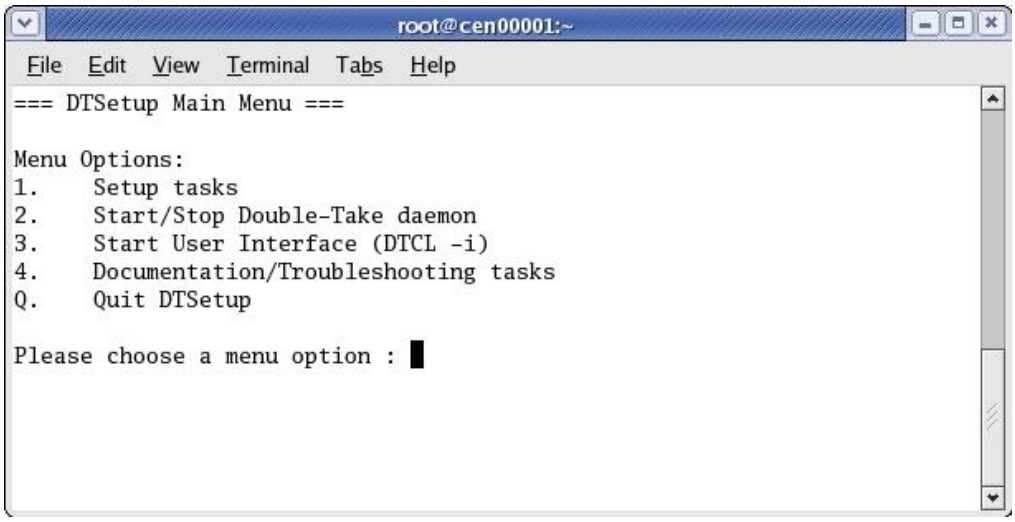
Running DTSetup

1. Run the DTSetup command from the shell prompt to start DTSetup. The command is case-sensitive.

   Do not run DTSetup using the sudo command. Use a real root shell to launch DTSetup instead, either by logging in as root on the console or by using the login session of a non-privileged user to run su - to start a root shell.

2. The first time you run DTSetup after an installation, you will be prompted to review the Vision Solutions license agreement. Review the agreement and accept the terms of agreement by typing yes. You cannot use Double-Take Availability without agreeing to the licensing terms.
3. When the DTSetup menu appears, enter the number of the menu option you want to access.



Setup tasks

The setup tasks are generally configured once. Some of the configuration options are per server, but the replication configuration options are for those servers in a source role. Select a link below to learn more about that setup task.

- Activation codes—These are the codes that license your Double-Take Availability servers.
- Security groups—Security groups provide access to Double-Take Availability.
- Replication configuration—You must configure each source server for replication. You have the choose of selecting file system replication, block device replication, or replicating an entire server.
- Server configuration—If desired, you can modify server settings through the Double-Take configuration file.
- Driver performance—If desired, you can specify Double-Take driver performance settings.

Activating your server

Before you can use Double-Take Availability, each source and target server must have a valid activation code, which is an alpha-numeric codes that applies the appropriate Double-Take Availability license to your installation.



Server activation can also be completed through the Replication Console. See Licensing a server.

1. Start DTSetup.
2. Select **Setup tasks**.
3. Select **Set Activation Code Menu**.
4. Select **Set Activation Code in /etc/DT/DT.conf**.
5. Enter your activation code and press Enter. The activation code will automatically be inserted into the configuration file. You are prompted to start the Double-Take service after the first installation, and you must restart the service each time the activation code is modified, such as after an upgrade.
6. Press Enter to return to the menu.
7. Press Q as many times as needed to return back to the main menu or to exit DTSetup.

Modifying security groups

During the installation, the user root is automatically added to the Double-Take Availability administrators security group. If you want to add other users or remove root, you will need to modify the security group configuration for each source and target

server. See Security for more details on the security groups and the privileges granted to each group.

1. Start DTSetup.
2. Select **Setup tasks**.
3. Select **Add/Remove users to Double-Take groups**.
4. Select the necessary menu options to add or remove groups to the administrator or monitors group as needed, and specify the user name when prompted.
5. When you have completed your security group modifications, press Q as many times as needed to return back to the main menu or to exit DTSetup.

Replication configuration

There are three types of replication configurations for your source server.

- File system replication—File system replication uses a Double-Take Availability file system mount point to capture data changes at the byte level. It is the most efficient type of replication. It provides a more detailed level of data selection, allowing you to include and exclude files and directories.
- Block device replication—Block device replication uses a loop device to capture data changes at the block level. Data selection is at the volume level. Generally, this replication configuration is ideal for non-file system type data, like raw databases.
- Full-server replication—Full-server replication captures the entire server for replication and protection using file system replication. When selecting this option, Double-Take Availability will automatically configure protection for all volumes. You can deselect any volume except root.

You cannot establish replication protection for the same set of data using more than one type of replication configuration. For example, /mydata cannot be protected by both file system replication and block device replication. However, you can protect two different file sets using two different replication configurations. For example /mydata1 and /mydata2 could be protected by file system replication and block device replication, respectively.

Configuring file system replication

In order to use Double-Take Availability replication, data that will be replicated must reside on a partition, which is specially mounted with the Double-Take Availability file system (DTFS) driver. DTFS is a transparent file system that monitors the storage file system so that data changes can be captured. Data on the file system can only be replicated if it is under DTFS mount points.

DTSetup allows you to configure entries in /etc/DT/dtfs_mounts to mount specific partitions as DTFS when Double-Take Availability is started.



When making replication configuration changes, stop any applications that may be running and restart them after the replication changes have been made. Double-Take Availability needs to be loaded on the file system before any applications, otherwise some data may not be replicated.

When protecting an NFS or Samba server, Double-Take Availability should be mounted just above the local file system where the exported NFS or Samba data resides. Additionally, NFS and/or Samba must be started after the Double-Take daemon.

DTFS mounts should remain mounted for the duration of the system uptime, if possible. Any file handles that a process has open on a DTFS mount location before DTFS is mounted will not be captured by replication. Since DTFS mounts are not automatically unmounted when the Double-Take daemon is stopped, the command service DTMount stop can be used to stop all DTFS mounts, if necessary.

1. Start DTSetup.
2. Select **Setup tasks**.
3. Select **Configure File System or Block Device Replication**.
4. Select **Configure File System Replication setup file**.
5. To add a DTFS mount, select **Add an entry to /etc/DT/dtfs_mounts**.
6. Specify the path of the highest subdirectory for which replication should be captured. This can be the path to a mount point or a subdirectory of a volume.

   Do not select the root as a replication mount point. If you want to include root as a replication mount point, you must configure full-server protection.

   If you have a single drive system, mount DTFS to any of the subdirectories directly under root.

7. You will be prompted whether you want to replicate access times. The access-time file property changes during a mirror and causes replication operations to be generated during a mirror and every time a file is read from or written to. This may have a negative effect on performance. If you do not replicate access times on both the source and target when the volume is initially mounted, you may increase the performance, especially during mirroring. (Access times from files closes are always replicated regardless of this configuration.) Specify yes or no to replicate access times.
8. To remove any DTFS mounts, select **Remove /Umount entries in /etc/DT/dtfs_mounts** and specify the path that you want to remove.

9. To immediately mount all of the entries in /etc/DT/dtfs_mounts, select **Mount all entries in /etc/DT/dtfs_mounts as DTFS**. If you do not select this option, you must stop and restart the daemon for the changes to take effect.
10. To immediately unmount all of the entries in /etc/DT/dtfs_mounts, select **Unmount DTFS from all entries in /etc/DT/dtfs_mounts**. This allows you to make changes to DTFS mounts without unloading the daemon.

    If you are using Konqueror, you can only unmount a replication mount point by rebooting the server.

11. When you have completed your file system replication configuration, press Q as many times as needed to return back to the main menu or to exit DTSetup.



You can also mount and unmount DTFS manually using the **Setup Tasks**, **Configure File System or Block Device Replication**, **Manual Replication Configuration menu** option. Changes made from this menu are not persisted between reboots/restarts.

Configuring block device replication

In order to use Double-Take Availability replication, data that will be replicated on a block device must be accessed through a loop device, which is specially attached using the Double-Take Availability loop driver (DTLOOP). DTLOOP allows the loop device to serve as an access point for operations performed on the block device so that data changes can be captured. Existing block devices may be available for replication, but the data on those block devices can only be replicated if they are accessed through the DTLOOP loop device. It is important that operations on the block device be made through the loop device only, or the operations will not be replicated. Failure to do so will result in corrupted data on the target.

DTSetup allows you to configure entries in /etc/DT/dtloop_devices to attach block devices as DTLOOP when Double-Take Availability is started.



If your block device being protected with DTLOOP has a file system on it, do not mount them from /etc/fstab. They should be mounted from an init script. DTMount must be started in the boot sequence before the script to mount the loop devices is executed in order to ensure that the loop devices have the replicated block devices associated with them. The script should then mount the loop device, not on the native block device.

When making replication configuration changes, stop any applications that may be running and restart them after the replication changes have been made.

Double-Take Availability needs to be loaded on the file system before any applications, otherwise some data may not be replicated.

> After the block device replication configuration is complete, applications must read and write through the /dev/loop# device in order for replication to work.

---

1. Start DTSetup.
2. Select **Setup tasks**.
3. Select **Configure File System or Block Device Replication**.
4. Select **Configure Block Device Replication setup file**.
5. If you want to see a list of block devices to which DTLOOP can be attached, select **List block devices on this system**. Swap and LVM physical partitions will not be included in the list.
6. Press Enter to continue.
7. To add a DTLOOP device, select **Add an entry to /etc/DT/dtloop_devices**.
8. Enter the path to the block device that is to be replicated and press Enter.
9. Enter the path to the loop device to use (/dev/loop#), if the same one should always be attached. If you use more than one loop device, you should assign a specific number to the loop device so it will persist beyond reboots. DTLOOP can also use the first one available, but that may mean it attaches to a different one on subsequent reboots/restarts, which may not be desirable. Press Enter to continue.
10. You will be asked if you want to attach at an offset into the block device and if you want to use an encrypted loop device. In general, these options can be left blank. See the losetup man page for more information on using encryption.
11. To remove any DTLOOP devices, select **Remove /Detach entries in /etc/DT/dtloop_devices** and specify the path that you want to remove.
12. To immediately attach all of the entries in /etc/DT/dtloop_devices, select **Attach all entries in /etc/DT/dtloop_devices to a loop device**. If you do not select this option, you must stop and restart the daemon for the changes to take effect.
13. To immediately detach all of the entries in /etc/DT/dtloop_devices, select **Detach loop devices from all entries in /etc/DT/dtloop_devices**. This allows you to make changes to DTLOOP without unloading the daemon.
14. When you have completed your block device replication configuration, press Q as many times as needed to return back to the main menu or to exit DTSetup.



You can also attach and detach DTLOOP manually using the **Setup Tasks**, **Configure File System or Block Device Replication**, **Manual Replication Configuration menu** option. Changes made from this menu are not persisted between reboots/restarts.

Configuring full-server protection

In order to use full-sever protection, Double-Take Availability must configure the root file system for replication. In order for the replication changes to be applied, you will have to reboot your server. Ideally, you will want to reboot the server immediately. However, if you have to schedule the reboot for a later time, mirroring will start immediately. Replication will not start until the server is rebooted. Without replication, data integrity cannot be maintained on the target.



When making replication configuration changes, stop any applications that may be running and restart them after the replication changes have been made. Double-Take Availability needs to be loaded on the file system before any applications, otherwise some data may not be replicated.

Full-server protection can also be completed through the Connection Wizard.

1. Start DTSetup.

   If the server has not been rebooted (forcing the Double-Take daemon to start), start the daemon manually.

2. Select **Setup tasks**.
3. Select **Configure File System or Block Device Replication**.
4. Select **Full Server Replication Configuration menu**.
5. Select **Add Full Server Protection**.
6. You will be prompted if you want to create the replication set. It can be created now or later through the Connection Wizard. Select yes or no.
7. If you chose to create the replication set now, enter a user that is a member of the Double-Take Availability administrators security group, and then specify the password associated with that user. Press Enter to continue.
8. If you want to view the contents of the replication set, select **View replication set contents**. After you have reviewed the contents, press Enter to continue.
9. If you need to remove full-server protection, select **Remove Full Server Protection**. Confirm that you want to remove full-server protection, and then press Enter to continue. A server reboot is required to fully remove full-server protection.
10. If you want to modify any of the data protected manually, you can select **Configure /etc/DT/dtfs_mounts manually** which will take you to the Configure File System Replication setup file menu.
11. When you have completed your full-server protection, press Q as many times as needed to return back to the main menu or to exit DTSetup.

Remember, in order for the replication changes to be applied, you will have to reboot your server. Ideally, you will want to reboot the server immediately. However, if you have to schedule the reboot for a later time, mirroring will start immediately. Replication will not start until the server is rebooted. Without replication, data integrity cannot be maintained on the target.

Configuring server settings

Server settings are available in various places. You can access them via the Replication Console, through DTCL See the Scripting Guide for details on accessing the server settings through DTCL. , or through DTSetup. Initially, the server settings file, /etc/DT/DT.conf, on the source and target is blank. To populate it with default values, start and stop the Double-Take daemon once.

1. Start DTSetup.
2. Select **Setup tasks**.
3. Select **Edit Double-Take config file**.
4. The server settings are listed in alphabetical order. Make modifications as necessary, using the control keys specified at the bottom of the page. For a complete list of each server setting, valid values, default values, and optional notes, see *Server Settings* in the *Scripting Guide*.
5. Press control-X to exit the configuration file.
6. Enter Yes or No to save any changes.
7. Press Q as many times as needed to return back to the main menu or to exit DTSetup.

Configuring driver performance settings

Driver settings provide configuration flexibility so you can adjust Double-Take Availability based on your servers, network, and replication requirements. You may want to modify driver settings on both the source and target.



Changing the driver performance settings can have a positive or negative impact on server performance. These settings are for advanced users. If you are uncertain how to best modify the driver performance settings, contact technical support.

1. Start DTSetup.
2. Select **Setup tasks**.
3. Select **Configure Double-Take driver performance**.
4. The current driver settings are displayed.
5. Select a driver setting to modify the option.

- **Toggle Adaptive Throttling**—You can toggle between enabling (true) and disabling (false) **Adaptive Throttling**. This occurs when kernel memory usage exceeds the **Throttling Start Level** percentage. When throttling is enabled, operations are delayed by, at most, the amount of time set in **Maximum Throttling Delay**, thus reducing kernel memory usage. Throttling stops when the kernel memory usage drops below the **Throttling Stop Level** percentage.
- **Toggle Forced Adaptive Throttling**—You can toggle between enabling (true) and disabling (false) **Forced Adaptive Throttling**. This causes all operations to be delayed by, at most, the amount of time in set in **Maximum Throttling Delay**, regardless of the kernel memory being used. **Adaptive Throttling** must be enabled (true) in order for **Forced Adaptive Throttling** to work.
- **Set Maximum Throttling Delay**—This option is the maximum time delay, in milliseconds, used by the driver during a system delay.
- **Set Throttling Delay Interval**—This option is the interval, in milliseconds, to check memory usage during a throttling delay. If a delay is no longer needed, the remainder of the delay is skipped.
- **Set Throttling Start Level**—Throttling starts when disk writes reach the specified percentage. This prevents the driver from stopping replication because memory has been exhausted.
- **Set Throttling Stop Level**—Throttling stops when disk writes reach the specified percentage.
- **Set Memory Usage Limit**— This option is the amount of kernel memory, in bytes, used for queuing replication operations. When this limit is exceeded, the driver will send an error to the daemon forcing a remirror of all active connections.
- **Set Maximum Write Buffer Size**— This option is the maximum amount of system memory, in bytes, allowed for a single write operation. Operations exceeding this amount are split into separate operations in the queue.

6. After you have completed your driver performance modifications, press Q as many times as needed to return back to the main menu or to exit DTSetup.

Starting and stopping the daemon

The Double-Take daemon will start automatically after Double-Take Availability is installed and the server is rebooted. You can start and stop the Double-Take daemon using this built-in DTSetup script.

1. Start DTSetup.
2. Select **Start/Stop Double-Take daemon**.
3. Select the necessary menu option to start or stop the daemon and handle the driver configuration.
   - **Start Double-Take and process driver config**—This option starts the Double-Take daemon and replication configuration is processed.
   - **Stop Double-Take but preserve driver config**—This option stops the Double-Take daemon but does not unload the drivers from the kernel or detach or unmount replicated devices.
   - **Restart service but preserve driver config**—This option does a full stop and start of the Double-Take daemon but does not unload the drivers from the kernel or detach or unmount replicated devices.
   - **Restart service and reset driver config**—This option does a full stop and start, completely unloading the daemon and drivers and then reloading them.
   - **Stop the running service and teardown driver config**—This option stops the Double-Take daemon and the drivers are unloaded from the kernel after detaching and unmounting replicated devices.
   - **Go to Replication Configuration menu**—This option takes you to **Setup Tasks**, **Configure File System or Block Device Replication**, **Manual Repliacation Configuration menu**. When you press Q to exit from that menu, you will return this menu.
4. When you have completed your starting and stopping tasks, press Q as many times as needed to return back to the main menu or to exit DTSetup.

Starting DTCL

You can launch the Double-Take Availability interactive command prompt which allows you to enter DTCL commands one at a time.

1. Start DTSetup.
2. Select **Start User Interface (DTCL -i)**.
3. Enter your DTCL commands one at a time at the **Command** prompt. For a complete list of DTCL commands, their syntax, and instructions for completing tasks using DTCL, see the *Scripting Guide*.
4. To exit the DTCL **Command** prompt, type exit.
5. When you have completed your DTCL tasks, press Q as many times as needed to return back to the main menu or to exit DTSetup.

Viewing documentation and troubleshooting tools

1. Start DTSetup.
2. Select **Documentation/Troubleshooting tasks**.
3. Select **View log files** to view the following log files. Double-Take Availability logs alerts, which are processing notifications, warnings, and error messages. The logs are written to disk.
   - **View /*.dtl in less**—This option uses the less file viewer program to view all of the Double-Take Availability logs, starting from the most recent.
   - **Follow the output of latest**—This option uses tail -f to watch the output of the Double-Take Availability logs in real-time.
   - **View /var/log/messages in less**—This option uses the less file viewer program to view the system log messages.
   - **Follow the output of /var/log/messages**—This option uses tail -f to watch the output of the system log messages in real-time.
4. Select one of the **Collect and package diagnostic info** selections to run the DTInfo script which collects configuration data. This can be useful when reporting problems to technical support. Depending on the diagnostic option you select, the amount of data to be collected varies between basic, detailed and full diagnostic information. You must have root (or uid 0 equivalent) to execute the diagnostics or to copy or read the resulting file.
5. Select **View user documentation** to view Double-Take Availability product documentation and several legal documents. DTSetup attempts to determine your viewers, although you can specify your viewer.
   - **View ReadMe HTML**—This option views the readme file which contains last minute release notes.
   - **View Scripting Guide HTML**—This option views the Scripting Guide which contains DTCL commands and scripting information.
   - **View User's Guide HTML**—This option views the User's Guide which contains instructions for using Double-Take Availability.
   - **View End User License Agreement TXT**—This option views the End User License Agreement legal document.
   - **View dtfs/dtrep module license TXT**—This option views the open source legal document.
   - **Change a document viewer**—This option allows you to specify a document viewer.
6. When you have completed your documentation and troubleshooting tasks, press Q as many times as needed to return back to the main menu or to exit DTSetup.

DTSetup menus

The following lists is an overview of the DTSetup menu system. Reference the links for complete details on completing tasks in DTSetup.

1. **Setup tasks**—Activation codes, security groups, replication configuration, server configuration, and driver performance settings. See Setup tasks.
    1. **Set Activation Code Menu**—See Activating your server.
    2. **Add/Remove users to Double-Take groups**—See Adding users to security groups.
    3. **Configure File System or Block Device Replication**—There are three types of replication configurations for your source server. See Replication configuration.
        1. **Configure File System Replication setup file**—See Configuring file system replication.
        2. **Configure Block Device Replication setup file**—See Configuring block device replication.
        3. **Manual Replication Configuration menu**—See Configuring file system replication and Configuring block device replication.
        4. **Full Server Replication Configuration menu**—See Configuring full-server replication.
    4. **Edit Double-Take config file**—See Configuring server settings.
    5. **Configure Double-Take driver performance**—See Configuring driver performance settings.
2. **Start/Stop Double-Take daemon**—See Starting and stopping the daemon.
3. **Start User Interface (DTCL -i)**—See Starting DTCL.
4. **Documentation/Troubleshooting tasks**—See Viewing documentation and troubleshooting tools.

Chapter 5 Clients

Double-Take Availability has two clients, the Replication Console and the Failover Control Center that control and manage your connections and failover. Both clients can be started from the Windows **Start** menu. You can also launch the Failover Control Center from the **Tools** menu in the Replication Console.

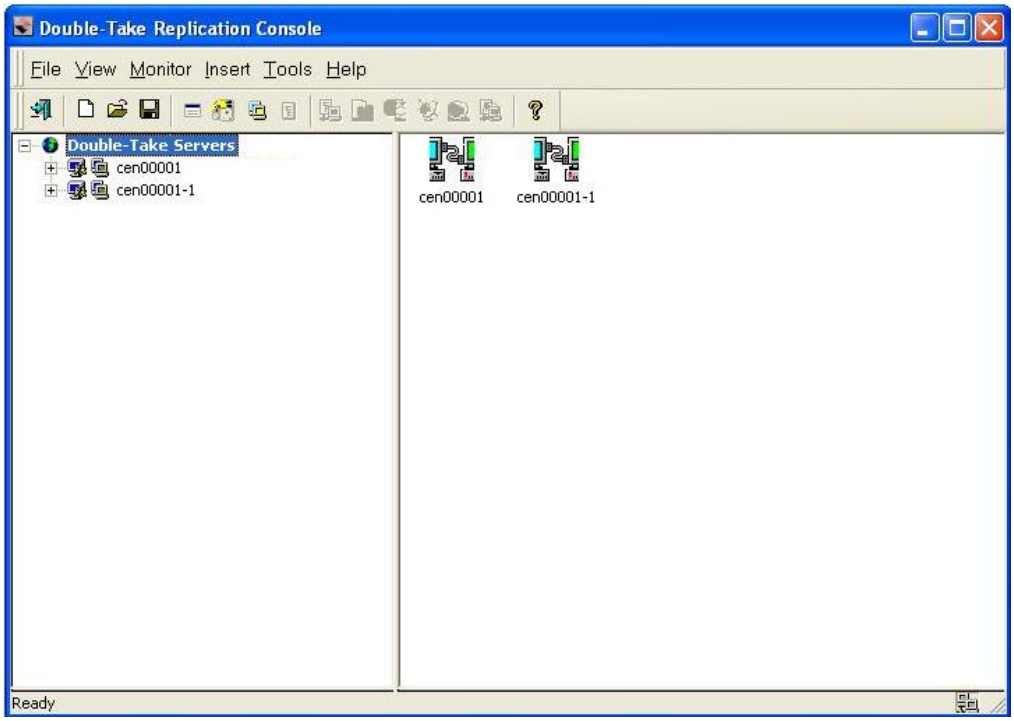


Double-Take Availability also has a scripting language which can be used in an interactive client or in scripts. For more information, see the *Scripting Guide*.

Replication Console

Start the Double-Take Availability Replication Console by selecting **Start**, **Programs**, **Double-Take for Linux**, **Double-Take Replication Console**.

From the Replication Console, you can manage, monitor, and control your Double-Take Availability 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.



Using Replication Console 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.

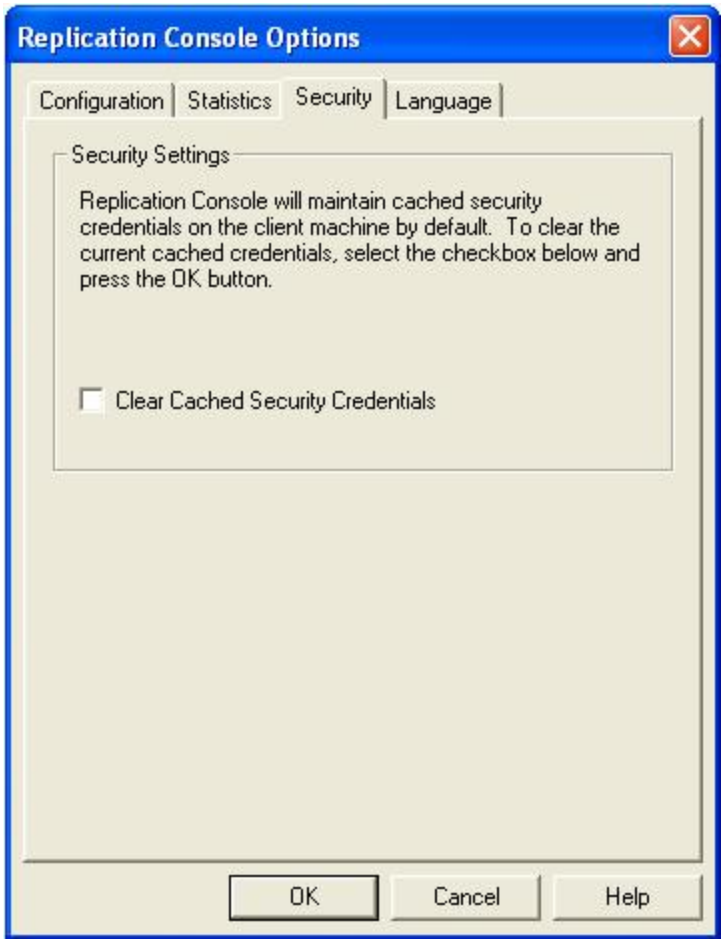
- **Saving a workspace**—As you size, add, or remove windows in the Replication Console, you can save the workspace to use later or use on another Double-Take Availability client machine. Select **File** and one of the following options.
- **Save Workspace**—Save the current workspace. If you have not 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 stored security credentials

Use the following steps to remove credentials cached in the Replication Console.

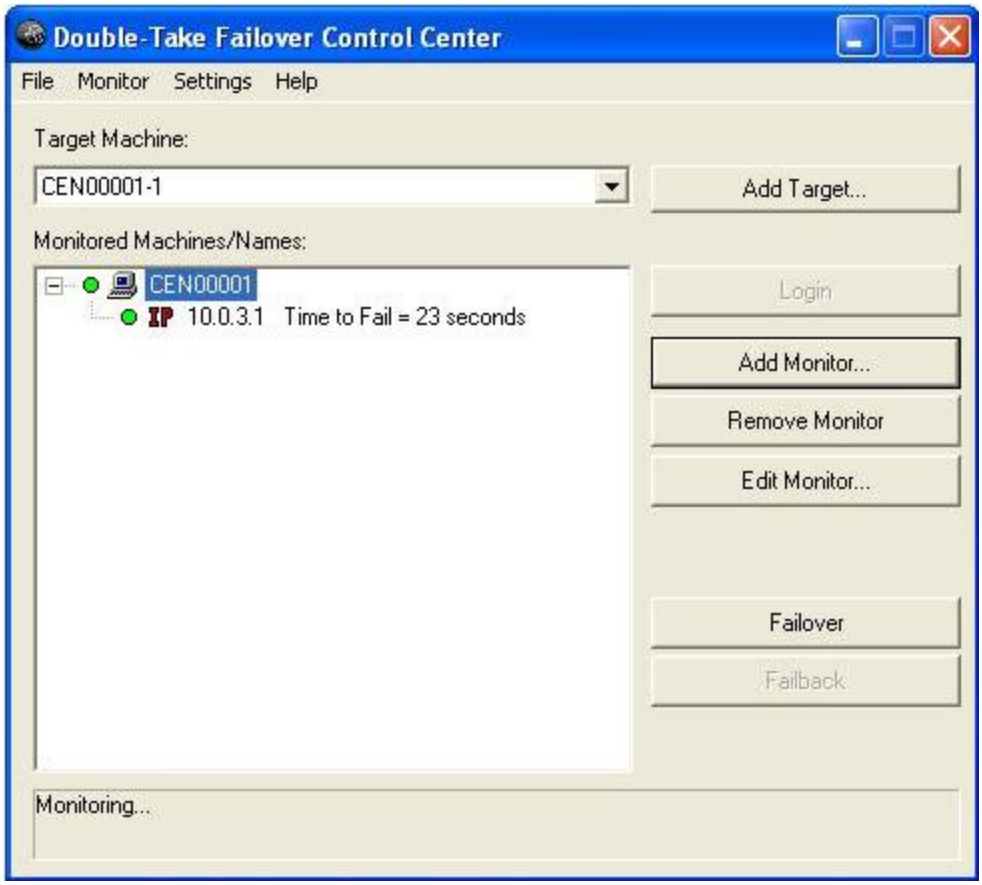
1. To access the credentials security option, select **File**, **Options** and select the **Security** tab.



2. To remove the security credentials, click **Clear Cached Security Credentials**.
3. Click **OK**.

Failover Control Center

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.



Setting the frequency of Failover Control Center console refreshes

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.

Chapter 6 Data protection

Protecting your 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.

Establishing a data 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. Start the Connection Wizard to establish your connection by selecting **Tools**, **Connection Wizard**.

   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.

   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 that you want to protect.

   Double-Take Availability will automatically attempt to log on to the selected source using previously cached credentials. If the logon is not successful, the Logon dialog box will appear prompting for your security identification.

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.

   Double-Take Availability will automatically attempt to log on to the selected target using previously cached credentials. If the logon is not successful, the Logon dialog box will appear prompting for your security identification.

6. Click **Next** to continue.
7. Select the type of connection you want to establish.
   - **Protect data**—This connection type protects only the data you will be selecting.
   - **Protect full server including system state and data**—This connection type protects both the source's system state, which is the configured operating system and installed applications, and all of the data on the server.

   Select **Protect data** and click **Next** to continue.
8. 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.
9. Click **Next** to continue.
10. 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.
11. 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_name.
    - **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.
    - **Custom**—To select a custom path, click once in the Target Path field and modify the drive and directory to the desired location.
12. Click **Next** to continue.

13. Review your selections on the summary screen. If your Connection Wizard settings are correct, establish your connection by completing one of the two options below.
    - 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.

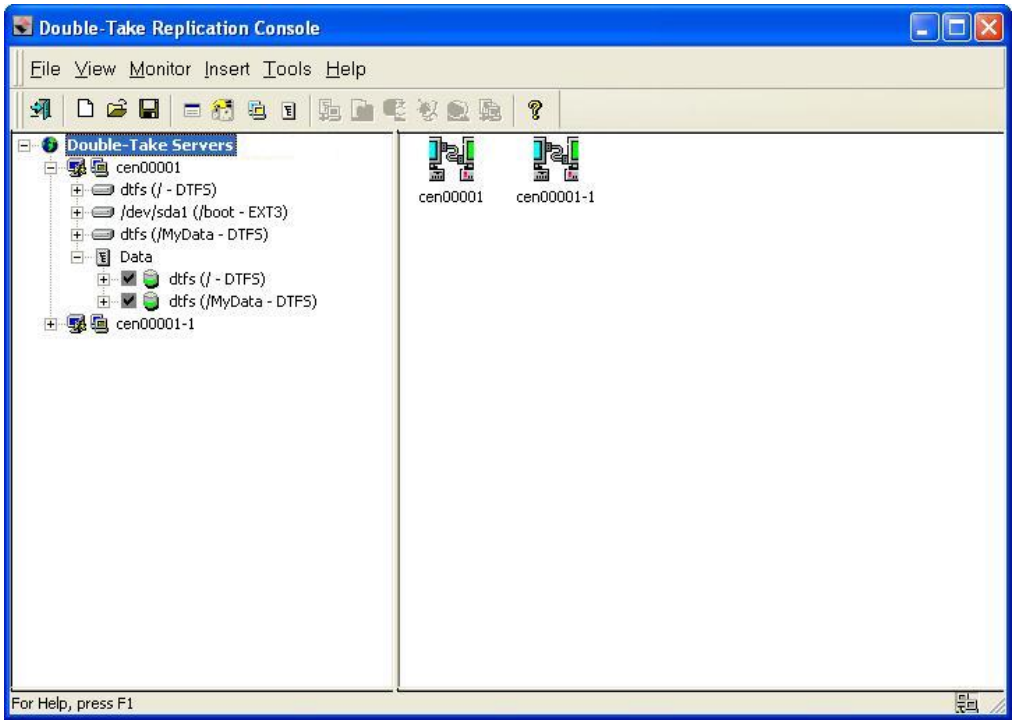
1. Highlight a source in the left pane of the Replication Console and select **Insert**, **Replication Set** from the menu bar. You can also right-click on the source name and select **New**, **Replication Set**.
2. A replication set icon appears in the left pane under the source. By default, it is named 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. The tree will display all mount points capable of supporting replication. Those that have been mounted for replication will be identified as DTFS mount points. If the mount point you want is not labeled as DTFS, you must mount it for replication through the DTSetup utility. Mounting can occur without restarting the daemon.



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, in a DTFS mount point, that you want to protect by selecting volumes, drives, directories, and/or specific files.



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

Replication sets should only include necessary data. Including data such as temporary files, logs, and/or locks will add unnecessary overhead and network traffic. For example, if you are using Samba, make sure that the location of the lock file (lock dir in samba.conf) is not a location in your Double-Take Availability replication set.

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.
6. If you need to select a block device for replication, right-click the replication set and select **Add Device**.
7. The block devices configured for Double-Take Availability replication are shown by default. Highlight the device to include in the replication set and click **OK**.

If the device you want to include is not displayed, you can click **Show Other Devices** to view all devices which are eligible for Double-Take Availability replication. You can select any of these devices, but you



cannot use them for Double-Take Availability replication until they are configured for Double-Take Availability replication. The status **no dtloop** indicates the device is not configured for Double-Take Availability replication.

Make sure your target has a partitioned device with sufficient space. It should be equal to or greater than the storage of the source device.

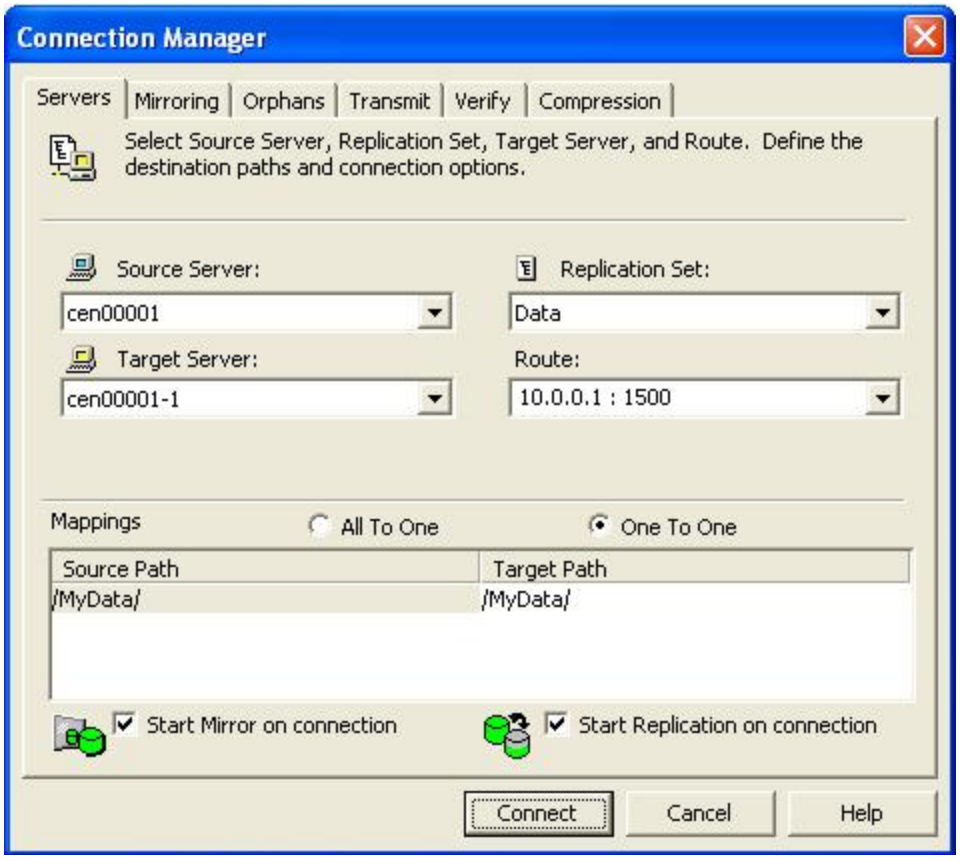
The partition size displayed may not match the output of the Linux df command. This is because df shows the size of the mounted file system not the underlying partition which may be larger. Additionally, Double-Take Availability uses powers of 1024 when computing GB, MB, and so on. The df command typically uses powers of 1000 and rounds up to the nearest whole value.

8. Repeat steps 6 and 7 for any additional devices.
9. Right-click the updated replication set icon and select **Save**.

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. 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 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. Select this option if you are protecting an entire server.
  - **One To One**—This option replicates data from the source to the same directory structure on the target. For example, /var/data and /usr/files on the source will be replicated to /var/data/ and /usr/files, respectively, on the target.

    Do not select the **One To One** option if you are protecting an entire server. You will overwrite essential files on your target causing it to become unstable or unusable.

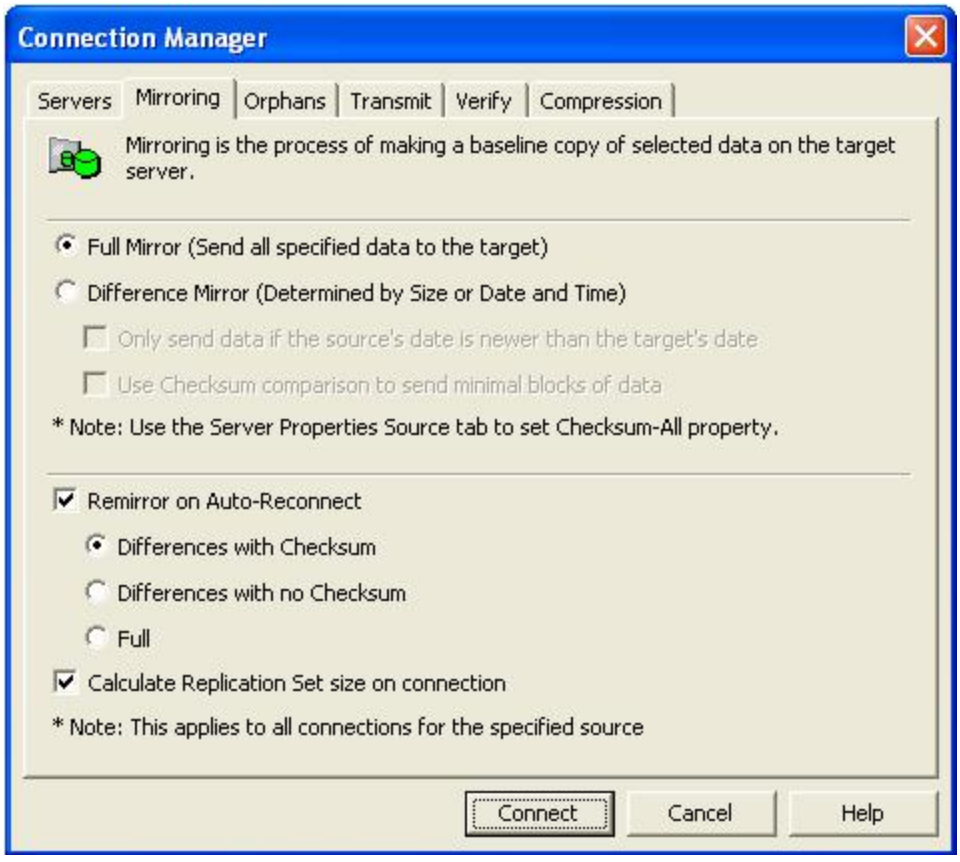
  - **Custom Location**—If the predefined options do not store the data in a location that is appropriate for your network operations, you can specify your own custom location where the replicated files will be sent. Click the target path and edit it, selecting the appropriate location.

- **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.

  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.

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.
- **Difference Mirror**—Only those files that are different based size or date and time (depending on files or block devices) will be sent from the source to the target.

- **Only send data if the source's date is newer than the target's date**—Only those files that are newer on the source are sent to the target.

  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 checksum comparison to send minimal blocks of data**—For those files flagged as different, the mirror performs a checksum comparison and only sends those blocks that are different.

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

    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.

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. Using the following table, locate and record your port settings for each of the Double-Take Availability ports. 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.

Replication Console Status Transmit Port

The Status Transmit Port sends and receives directed UDP communications to display status and statistics in the Replication Console. The default setting is 1505.

Replication Console Heartbeat Advertisement

The Heartbeat Advertisement port sends and receives broadcast UDP communications to populate the Replication Console tree with Double-Take Availability servers. The default setting is 1500.

Failover Control Center Service Transmit Port

The Service Transmit Port sends and receives TCP communication between Double-Take Availability servers

and between Double-Take Availability servers and Double-Take Availability clients. The default setting is 1500.

Failover Control Center Heartbeat Listen Port

The Heartbeat Listen Port send and receives broadcast UDP communications to populate the Failover Control Center with Double-Take Availability servers. The default setting is 1500.

Double-Take Availability Server Service Listen Port

The Service Listen Port sends and receives TCP communication between Double-Take Availability servers and between Double-Take Availability servers and Double-Take Availability clients. The default setting is 1500.

Double-Take Availability Server Heartbeat Transmit Port

The Heartbeat Advertisement port sends and receives broadcast UDP communications to populate the Replication Console tree with Double-Take Availability servers. The default setting is 1500.

Double-Take Availability Server Status Listen Port

The Status Listen Port sends directed UDP communications to display status and statistics in the Replication Console. The default setting is 1505.

Double-Take Availability Server Statistics Logging Port

The port used for statistics logging is not available through a client. You must use the get and set DTCL commands to modify that port. See the Scripting Guide for details on the commands and the StatsPort option. The default setting is 1506.

---

2. You need to configure your hardware so that Double-Take Availability traffic is permitted access through the router and directed appropriately. Using the port information from the previous section, 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.

- In addition to UDP heartbeats, 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. If your network is configured to propagate UDP broadcasts, your servers will be populated in the Replication Console from across the router. If not, you have to 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, you can use the Connection Wizard or the Connection Manager to establish your connection.

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.

1. Before and after simulating your connection, you should gather network and system information specific to Double-Take Availability operations. Use DTSetup to run DTInfo to automatically collect this data.
2. Select the DTSetup option for troubleshooting, then select the option for basic diagnostics.
3. When you run the diagnostics, it may take several minutes for it to finish processing. When it is complete, a .tar.gz file will be created in /var/run/etc/DT/. The file name will have DTInfo with the date and time. You must have root (or uid 0 equivalent) to execute the diagnostics or to copy or read the resulting file.
4. Opening 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.
5. 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.
6. Click **Connect** to establish the connection. The simulation data will be logged to the Double-Take Availability statistics file.
7. Repeat steps 1-3 to run the diagnostics utility after the simulation is complete.

Chapter 7 Establishing full-server protection

To establish full-server protection, you will use the Connection Wizard which will guide you through the process of protecting your entire server. It helps you select a source and target, and then it automatically creates the replication set for you.

1. Start the Connection Wizard to establish your connection by selecting **Tools**, **Connection Wizard**.

   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.

   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 that you want to protect.

   Double-Take Availability will automatically attempt to log on to the selected source using previously cached credentials. If the logon is not successful, the Logon dialog box will appear prompting for your security identification.

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.

   Double-Take Availability will automatically attempt to log on to the selected target using previously cached credentials. If the logon is not successful, the Logon dialog box will appear prompting for your security identification.

6. Click **Next** to continue.
7. Select the type of connection you want to establish.
    - **Protect data**—This connection type protects only the data you will be selecting.
    - **Protect full server including system state and data**—This connection type protects both the source's system state, which is the configured operating system and installed applications, and all of the data on the server.

    Select **Protect full server including system state and data**. Once you proceed, Double-Take Availability will automatically configure the root file system for replication. In order for the replication changes to be applied, you will have to reboot your server. Ideally, you will want to reboot the server immediately. However, if you have to schedule the reboot for a later time, mirroring will start immediately. Replication will not start until the server is rebooted. Without replication, data integrity cannot be maintained on the target.

    If you need to remove the full server protection configuration, you must use DTSetup.
8. Click **Next** to continue.
9. Reboot your source server now if possible. After the source has been rebooted, acknowledge the reboot prompt. Remember, you must reboot your source to start replication.
10. By default, all volumes are automatically selected for protection. The protection rules have been designed to accurately select the necessary files to complete full-server protection. You should not have to make modifications. If required, you can deselect volumes that you do not want to protect by clearing the checkbox. If any deselected volumes contain application data, the integrity of your installed applications will not be maintained. You cannot clear the checkbox for the root volume because that is where the system state information is stored. A replication set will automatically be created using the volumes you select. Click **Next** to continue.
11. Select a location on the target to store the source image and data. In order to distinguish this image from images of other servers, the source name will automatically be appended to the path you select and will be displayed on-screen. 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 two options below.
    - If you do not want to set advanced options, click **Finish**. The Connection Wizard will close and the connection will be established. Mirroring will begin,

and if the source has been rebooted prior to the connection being established, replication will also 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.

Chapter 8 Protection monitoring

Double-Take Availability flexibility offers a wide variety of methods for monitoring your protection.

- Monitoring a connection through the Replication Console—You can use the build-in monitoring in the Replication Console to watch statistics and see at-a-glance the health of your protection.
- Log files—The Double-Take Availability log files provide notification, warning, and error processing messages.
- Monitoring the Linux system log—Double-Take Availability generates Linux system messages.
- Statistics—Additional statistics are available outside of the Replication Console.
- SNMP—Both statistics and processing messages are available through SNMP.

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.

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.

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.

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.

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.



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 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.

—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.

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.

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.

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

The following topics are available for the log file.

- Viewing the log files through a text editor
- Viewing the Double-Take Availability log file through the Replication Console
- Configuring the properties of theDouble-Take Availability log file
- Double-Take Availability log messages

Viewing the log files 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
```

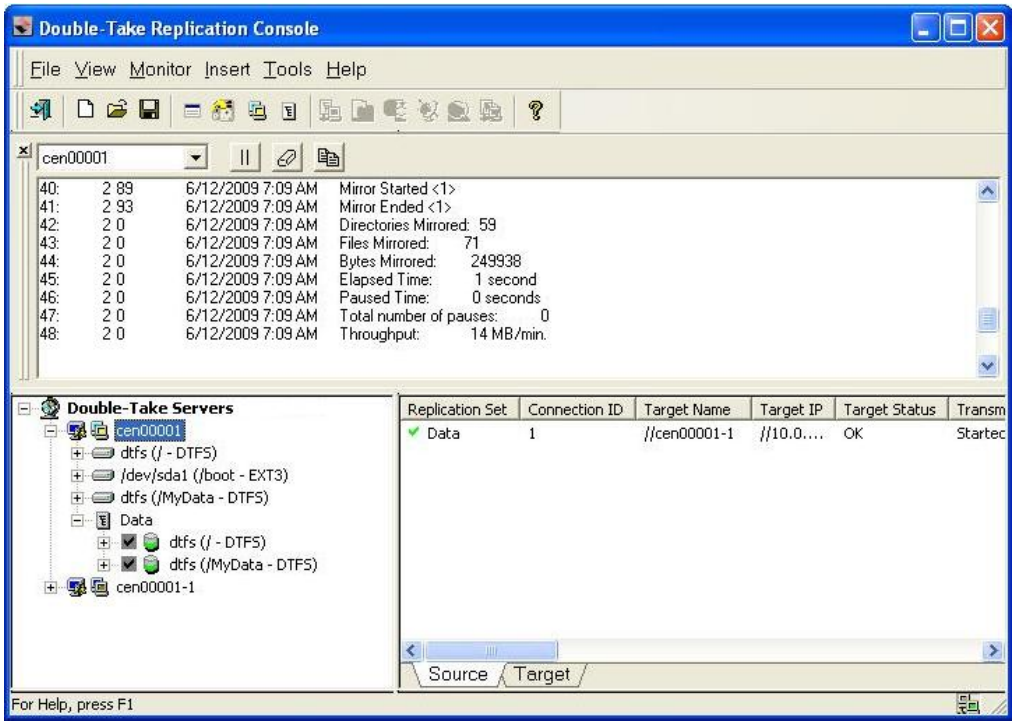
Sample Replication Console log file

```
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\Vision
Solutions\Double-Take for Linux\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).
```

Viewing the Double-Take Availability log file through the Replication Console

In addition to the statistics and status shown in the Replication Console, you can also open a message window to view the Double-Take Availability log file.

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.



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.

Configuring the properties of the Double-Take Availability 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 1048576 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.

   If you change the **Maximum Length** or **Maximum Files**, you must restart the Double-Take daemon 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. The ID appears in column 7 of the log file, and the message appears in column 8.



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 the Linux system log

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 the Linux system log. The location of the log file depends on the configuration of /etc/syslog.conf, however, by default, it is /var/log/messages. The following table identifies the events generated by Double-Take Availability.

---

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.

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.

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.

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.

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.

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.

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.).

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.

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.

E-mailing system messages

You can e-mail system messages to specified 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. To enable e-mail notification for a server, right-click the server in the left pane of the Replication Console and select **Properties**.
2. Select the **E-mail Notification** tab.
3. Select **Enable notification**.

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

4. Specify your e-mail settings.
   - **Mail Server (SMTP)**—Specify the name of your SMTP mail server.

     Specifying an SMTP server is the preferred method because it provides a direct connection between the mail server and Double-Take Availability, which decreases message latency and allows for better logging when the mail server cannot be reached.

     If you do not specify an SMTP server, Double-Take Availability will attempt to use the Linux mail command. The success will depend on how the local mail system is configured. Double-Take Availability will be able to reach any address that the mail command can reach.

   - **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.



You can test e-mail notification by specifying the options on the E-mail Notification tab and clicking **Test**. 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.

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 daemon, 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.

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.



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

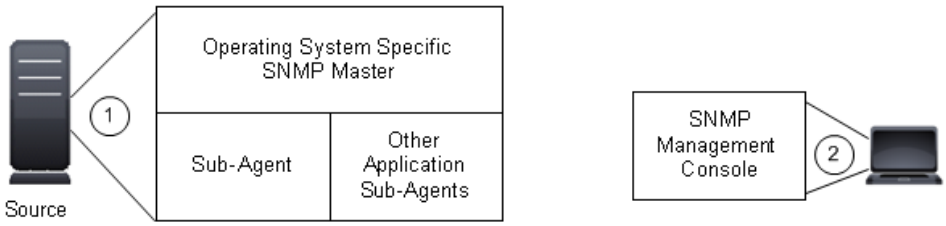
---

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 monitors Double-Take Availability and can be managed from an SNMP Management Console.

Double-Take Availability installs two components to work with SNMP.

1. 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.
2. 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 NSI-DT.mib and meets SNMP standards.



The Double-Take Availability SNMP sub-agent is only supported for NET - SNMP v2c.

Configuring SNMP on your server

SNMP must be installed, configured, and working on your server.

1. Stop the SNMP daemon (snmpd).
2. Make a backup copy of the SNMP configuration file snmpd.conf.
3. In order to run the Double-Take Availability sub-agent, master agentx is needed to turn on agentx. Add the master agent line to snmpd.conf.

   master agentx
4. So that the SNMP daemon can locate the Double-Take Availability MIB, add the path to the Double-Take Availability MIB by adding an entry to the snmp.conf file.

   mibfile /usr/share/snmp/mibs/NSI-DT-MIB.txt
5. Restart the SNMP daemon (snmpd).

6. Start the master agent.

   #snmpd -f -Le -x /var/agentx/master &

7. Start the Double-Take Availability SNMP sub-agent.

   #DTSubAgent >& /dev/null &

   Instead of starting the master agent and Double-Take Availability sub-agent separately, you can add them both to init.d to start them automatically.

8. You can test SNMP by trying either of the following commands.

   #snmpget -v2c -c public localhost dtGeneral.dtUpTime.0

   #snmpget -v2c -c public localhost NSI-MIB::dtUpTime

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 from the source to a target

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

---

Chapter 9 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 daemon is restarted. For example, if the Double-Take daemon 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 daemon.

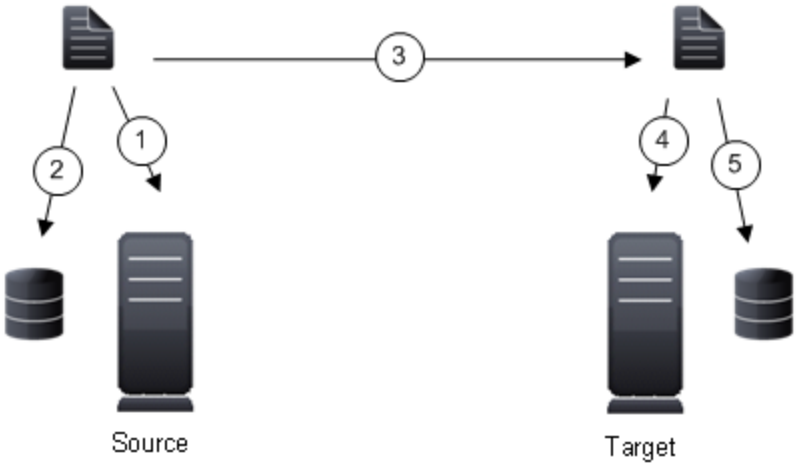
You can perform the following functions to manage your connections.

- Queuing Double-Take Availability data
- Reconnecting automatically
- Pausing and resuming target processing
- 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 backup.

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.



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.

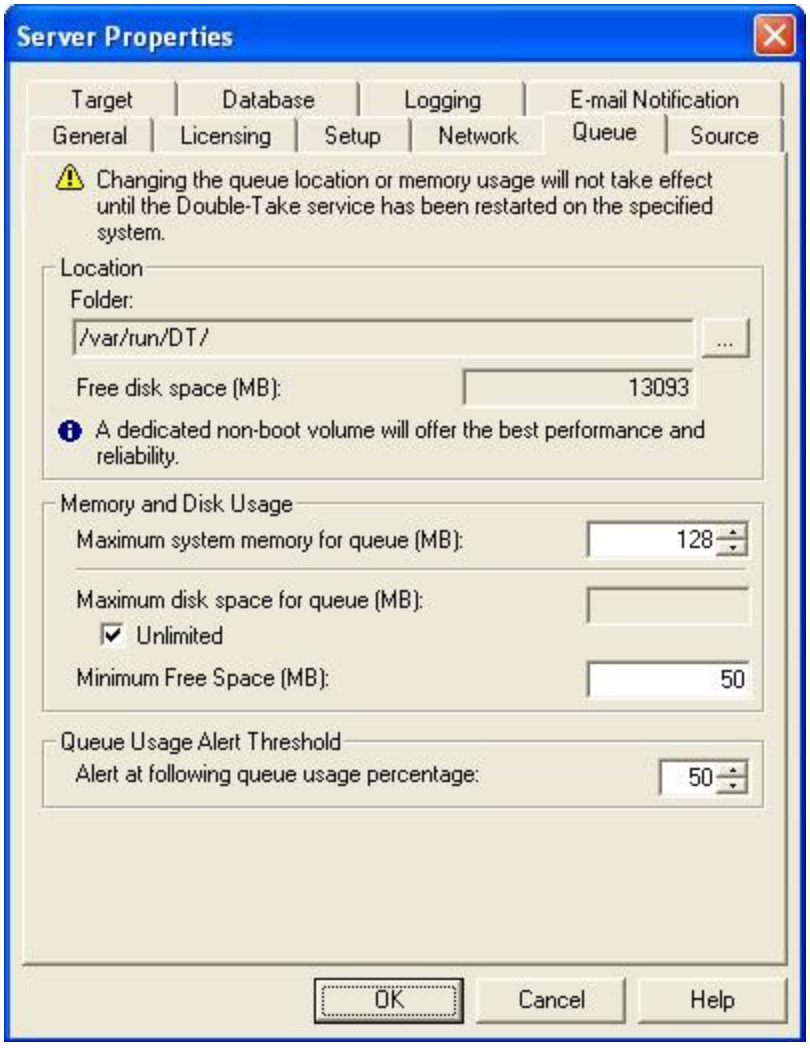
3. When system memory is full, the most recent changed data is added to the disk queue, as described in step 2. This means that system memory contains the oldest data. Therefore, when data is transmitted to the target, Double-Take Availability pulls the data from system memory and sends it. This ensures 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. 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 daemon has been restarted on the server.

  Select a location on a volume that will have minimal impact on the operating system and applications being protected. For best results and reliability, this

should be a dedicated, non-boot volume. The disk queue should not be on the same physical or logical volume as the data being replicated.



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 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 of memory is used. 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 daemon 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**.



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.

- Specify the **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 daemon shutdown**—If the target daemon is stopped and restarted, there could have been data in the target queue when the daemon was stopped.

To prevent any loss of data, the Double-Take daemon 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 daemon is stopped. If Double-Take Availability is able to successfully persist this information, when the Double-Take daemon 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 daemon 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.



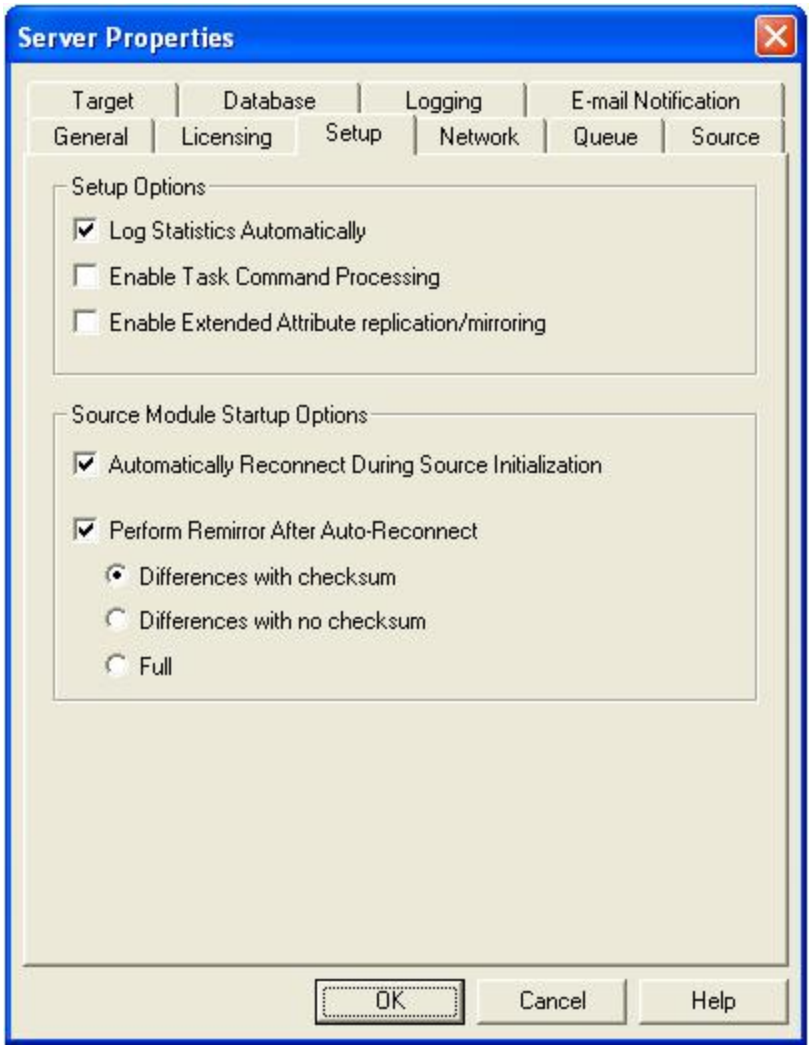
If you are experiencing frequent auto-disconnects, you may want to increase the amount of disk space on the volume where the Double-Take Availability queue is located or move the disk queue to a larger volume.

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. Right-click the source server on the left pane of the Replication Console and select **Properties**.
2. 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 may want to pause the target while you perform a backup of the target data, 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, right-click a target server on the left pane of the Replication Console and 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**.



If you have multiple connections to the same target, all connections will be paused and resumed.

Disconnecting a connection

To disconnect a Double-Take Availability connection, right-click the connection on the right pane of the Replication Console and select **Disconnect**. The source and target will be disconnected.



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.

Chapter 10 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
- 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. Right-click the connection on the right pane of the Replication Console and 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.

    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

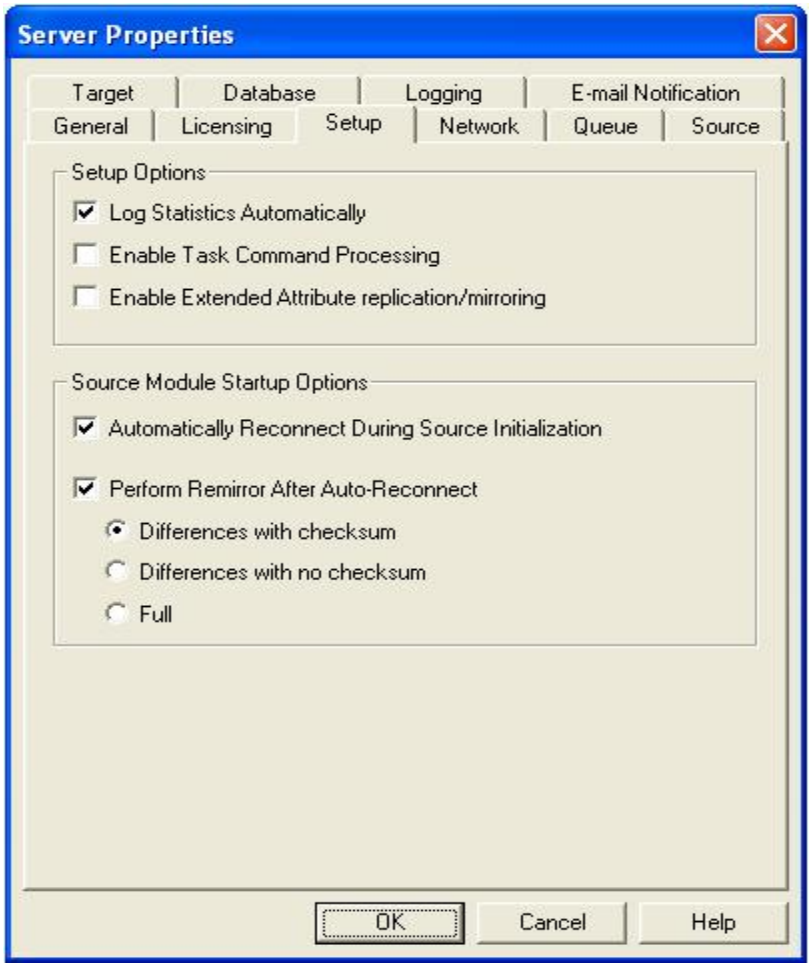
- **File Differences**—Any file that is different on the source and target based on the date, time, and/or size is transmitted to the target. The mirror sends the entire file.
- **File Differences and Only if Source is Newer**—Any file that is newer on the source than on the target based on date and/or 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.

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. Right-click a server in the left pane of the Replication Console and select **Properties**.
2. Select the **Setup** tab.



3. Verify that the **Perform Remirror After Auto-Reconnect** check box is selected to initiate an auto-remirror after an auto-reconnect.



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



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.

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

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.



Orphan file configuration is a per target option. All connections to the same target will have the same orphan file configuration.

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, carefully review any replication set rules that use wildcard definitions. If you have specified wildcards to be excluded from your replication set, files matching those wildcards will also be excluded from orphan file processing and will not be deleted from the target. However, if you have specified wildcards to be included in your replication, those files that fall outside the wildcard inclusion rule will be considered orphans and will be deleted from the target.

1. If you want to preview which files are identified as orphan files, right-click an established connection and select **Remove Orphans**, **Preview**. Check the log file on the target for the list of orphaned files.
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.



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.

Chapter 11 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 all file and directory data in the supported Linux file systems. Double-Take Availability does not replicate items that are not stored on the file system, such as pseudo-file systems like /proc and /sys. In addition, note the following.

- Data that will be replicated must reside on a Double-Take Availability partition, which is specially mounted with the Double-Take Availability file system (DTFS) driver. DTFS is a virtual file system that monitors the storage file system so that data changes can be captured. DTFS is a transparent file system. Existing mount points may be available for replication, but the data on those mount points can only be replicated if they are DTFS mount points. The DTSetup utility allows you to specify which mount points should be DTFS mount points.
- DTFS mount points are identified in the Replication Console by dtfs in the mount point name.
- Double-Take Availability is compatible with NFS and Samba services as long as they are mounted on top of Double-Take Availability. Double-Take Availability should be mounted just above the storage file system. Additionally, NFS and Samba should be started after the Double-Take daemon.
- Do not select the root as a replication mount point.
- If you have a single drive system, mount DTFS to any of the subdirectories directly under root that you want to replicate.
- If you select data stored on a recursive mount point for replication, a mirror will never finish. Double-Take Availability does not check for data stored on recursive mount points.
- If any directory or file contained in your replication set specifically denies permission to the account running the Double-Take daemon, the attributes of the file on the target will not be updated because of the lack of access.
- If you are using soft links, keep in mind the following.
  - If a soft link to a directory is part of a replication set rule's path above the entry point to the replication set data, that link will be created on the target as a regular directory if it must be created as part of the target path.
  - If a soft link exists in a replication set (or is moved into a replication set) and points to a file or directory inside the replication set, Double-Take Availability will remap the path contained in that link based on the Double-Take Availability target path when the option RemapLink is set to the default value (1). If RemapLink is set to zero (0), the path contained in the link will retain its original mapping.
  - If a soft link exists in a replication set (or is moved into a replication set) and points to a file or directory outside the replication set, the path contained in that link will retain its original mapping and is not affected by the RemapLink option.

- If a soft link is moved out of or deleted from a replication set on the source, that link will be deleted from the target.
- If a soft link to a file is copied into a replication set on the source and the operating system copies the file that the link pointed to rather than the link itself, then Double-Take Availability replicates the file copied by the operating system to the target. If the operating system does not follow the link, only the link is copied.
- If a soft link to a directory is copied into a replication set on the source and the operating system copies the directory and all of its contents that the link pointed to rather than the link itself, then Double-Take Availability replicates the directory and its contents copied by the operating system to the target. If the operating system does not follow the link, only the link is copied.
- If any operating system commands, such as chmod or chown, is directed at a soft link on the source and the operating system redirects the action to the file or directory which the link references, then if the file or directory referenced by the link is in a replication set, the operation will be replicated for that file to the target.
- The operating system redirects all writes to soft links to the file referenced by the link. Therefore, if the file referenced by the symbolic link is in a replication set, the write operation will be replicated to the target.

- If you are using hard links, keep in mind the following.
  - If a hard link exists (or is created) only inside the replication set on the source, having no locations outside the replication set, the linked file will be mirrored to the target for all locations and those locations will be linked if all link locations on the target exist on the same partition.
  - If a hard link crosses the boundaries of a replication set on the source, having locations both inside and outside the replication set, the linked file will be mirrored to the target for only those locations inside the replication set on the source, and those locations will be linked on the target if all link locations exist on the same partition.
  - If a hard link is created on the source linking a file outside the replication set to a location inside the replication set, the linked file will be created on the target in the location defined by the link inside the replication set and will be linked to any other locations for that file which exist inside the replication set.
  - If any hard link location is moved from outside the replication set into the replication set on the source, the link will not be replicated to the target even if other link locations already exist inside the replication set, but the linked file will be created on the target in the location defined by the link.
  - If any hard link location existing inside the replication set is moved within the replication set on the source, the move will be replicated to the target and the link will be maintained if the new link location does not cross partitions in the target path.

- If any hard link location existing inside the replication set is moved out of the replication set, that file or linked location will be deleted on the target.
- If a hard linked file is copied from any location inside or outside the replication set to a location inside the replication set on the source, the copy will be replicated to the target.
- If a hard linked file has a location in the replication set and any of the operating system commands, such as chmod or chown, are directed at that file from a location inside the replication set, the modification to the file will be replicated to the target. Operations on hard links outside of the replication set are not replicated.
- If a hard linked file has a location in the replication set and a write operation is directed at that file from inside the replication set, the write operation will be replicated to the target. Operations on hard links outside of the replication set are not replicated.
- If any hard link location existing inside the replication set is deleted on the source, that file or linked location will be deleted from the target.
- If a hard link location is created on the same volume but outside of the path where DTFS is mounted, the hard link operation will fail because the operating system regards the DTFS mount as a separate volume.

- If you are using nested mount points, keep in mind that DTFS does not traverse mount point boundaries, so DTFS must be mounted to each volume in the hierarchy in the proper order. You must manage the sequence of the combined native and DTFS mount operations either manually or by modifying /etc/fstab to mount the file systems in the correct order. For example, the following entries in /etc/fstab ensure the native volumes and their respective DTFS mounts will occur in the proper sequence for this configuration.

/dev/sda1	/level1	ext3	noatime	0	0
dtfs	/level1	dtfs	norepatime	0	0
/dev/sda2	/level1/level2	ext3	noatime	0	0
dtfs	/level1/level2	dtfs	norepatime	0	0
/dev/sda3	/level1/level2/level3	ext3	noatime	0	0
dtfs	/level1/level2/level3	dtfs	norepatime	0	0

If you want to make sure all DTFS mount points are accounted for, view the list from the start and stop daemon menu in DTSetup or by doing a cat of /etc/DT/dtfs_mounts.

When managing these mount points manually, the lower or subordinate native mount points must not be mounted until DTFS is mounted for the parent mount.

Using the same example and starting with all native volumes mounted, the following steps would be required.

1. umount /level1/level2/level3
2. umount /level1/level2/
3. mount dtfs –t dtfs –o norepatime /level1
4. mount /dev/sda2 –t ext3 –o noatime /level1/level2
5. mount dtfs –t dtfs –o norepatime /level1/level2
6. mount /dev/sda3 –t ext3 –o noatime /level1/level2/level3
7. mount dtfs –t dtfs –o norepatime /level1/level2/level3

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 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:

- **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 not to be included in 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.

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 4094 characters. Paths longer than 4094 characters will be skipped and logged to the Double-Take Availability log file and the Linux system log.

- 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 /var/data included recursively and /var/data/old included nonrecursively, the top level rule, /var/data/, will take precedence and the rule /var/data/old will be discarded. If the rules are different types (for example, /var/data is included and /var/data/old is excluded), both rules will be applied as specified.

- **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.

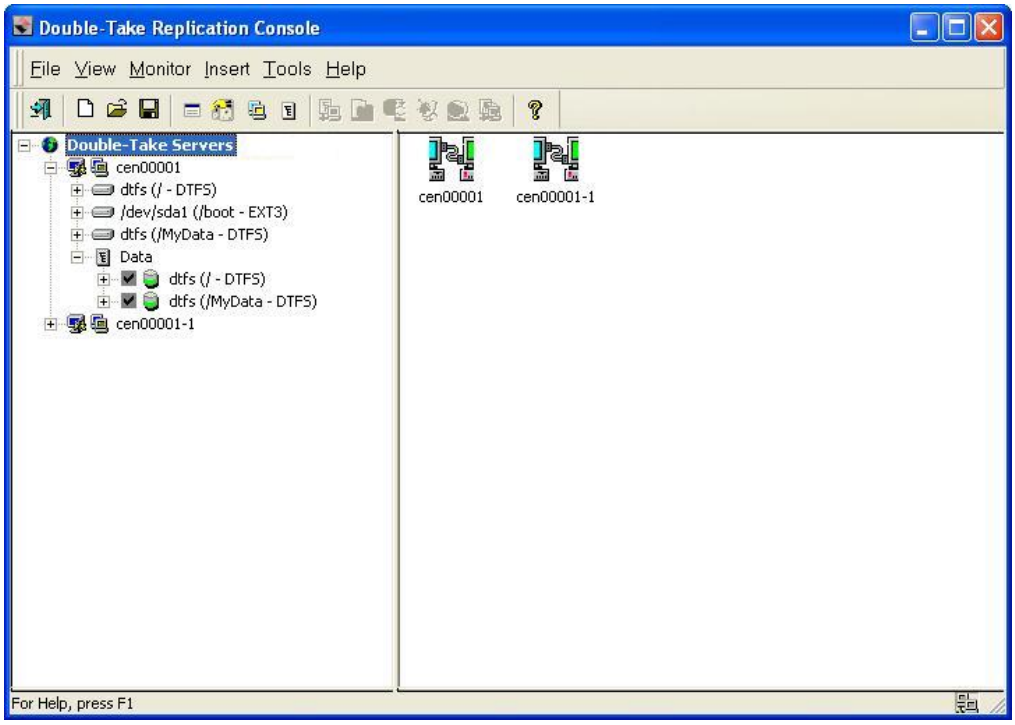
1. Highlight a source in the left pane of the Replication Console and select **Insert**, **Replication Set** from the menu bar. You can also right-click on the source name and select **New**, **Replication Set**.
2. A replication set icon appears in the left pane under the source. By default, it is named 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. The tree will display all mount points capable of supporting replication. Those that have been mounted for replication will be identified as DTFS mount points. If the mount point you want is not labeled as DTFS, you must mount it for replication through the DTSetup utility. Mounting can occur without restarting the daemon.



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, in a DTFS mount point, that you want to protect by selecting volumes, drives, directories, and/or specific files.



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

Replication sets should only include necessary data. Including data such as temporary files, logs, and/or locks will add unnecessary overhead and network traffic. For example, if you are using Samba, make sure that the location of the lock file (lock dir in samba.conf) is not a location in your Double-Take Availability replication set.

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.
6. If you need to select a block device for replication, right-click the replication set and select **Add Device**.
7. The block devices configured for Double-Take Availability replication are shown by default. Highlight the device to include in the replication set and click **OK**.

If the device you want to include is not displayed, you can click **Show Other Devices** to view all devices which are eligible for Double-Take Availability replication. You can select any of these devices, but you



cannot use them for Double-Take Availability replication until they are configured for Double-Take Availability replication. The status **no dtloop** indicates the device is not configured for Double-Take Availability replication.

Make sure your target has a partitioned device with sufficient space. It should be equal to or greater than the storage of the source device.

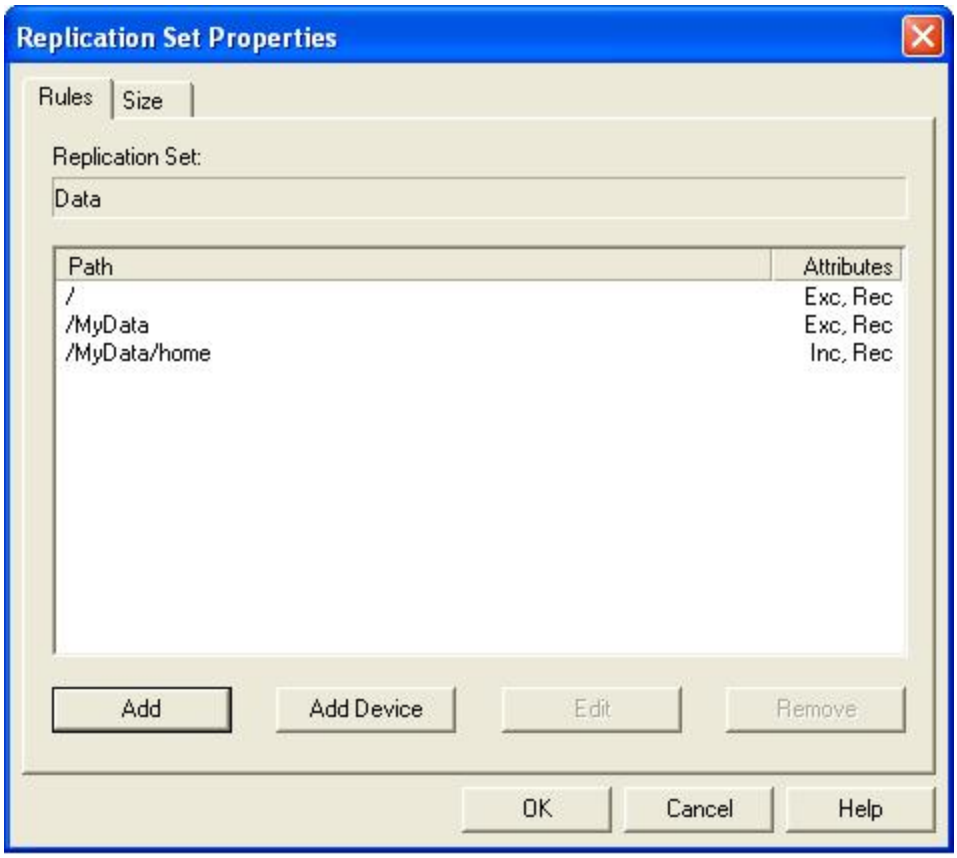
The partition size displayed may not match the output of the Linux df command. This is because df shows the size of the mounted file system not the underlying partition which may be larger. Additionally, Double-Take Availability uses powers of 1024 when computing GB, MB, and so on. The df command typically uses powers of 1000 and rounds up to the nearest whole value.

8. Repeat steps 6 and 7 for any additional devices.
9. Right-click the updated replication set icon and select **Save**.

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. 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.
2. 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.

3. From the Replication Set Properties dialog box, click **Add**.
4. 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**.
5. If you need to select block devices for replication, click **Add Device**. The block devices configured for Double-Take Availability replication are shown by default. Highlight the device to include in the replication set and click **OK**. If the device you want to include is not displayed, you can click **Show Other Devices** to view all devices which are eligible for Double-Take Availability replication. You can select any of these devices, but you cannot use them for Double-Take Availability replication until they are configured for Double-Take Availability replication. The status no dtloop indicates the device is not configured for Double-Take Availability replication.
6. If you need to edit an existing rule, highlight it and click **Edit**.

7. If you need to remove a rule, highlight it and click **Remove**.
8. 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.
9. Right-click the replication set icon and select **Save**. A saved replication set icon will change from red to black.

Selecting a block device for replication

Double-Take Availability allows you to select block devices for replication.

1. In the left pane, right-click the replication set that should include the block device and select **Add Device**.
2. The block devices configured for Double-Take Availability replication are shown by default. Highlight the device to include in the replication set and click **OK**.

   If the device you want to include is not displayed, you can click **Show Other Devices** to view all devices which are eligible for Double-Take Availability replication. You can select any of these devices, but you cannot use them for Double-Take Availability replication until they are configured for Double-Take Availability replication. The status **no dtloop** indicates the device is not configured for Double-Take Availability replication.

   Make sure your target has a partitioned device with sufficient space. It should be equal to or greater than the storage of the source device.

   The partition size displayed may not match the output of the Linux df command. This is because df shows the size of the mounted file system not the underlying partition which may be larger. Additionally, Double-Take Availability uses powers of 1024 when computing GB, MB, and so on. The df command typically uses powers of 1000 and rounds up to the nearest whole value.

3. Repeat steps 1 and 2 for any additional devices.

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. In the left pane, highlight the replication set you want to modify and expand the volume and directory levels as needed.
2. 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.
3. 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.



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.

Renaming and copying a replication set

To rename or copy a replication set, 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.

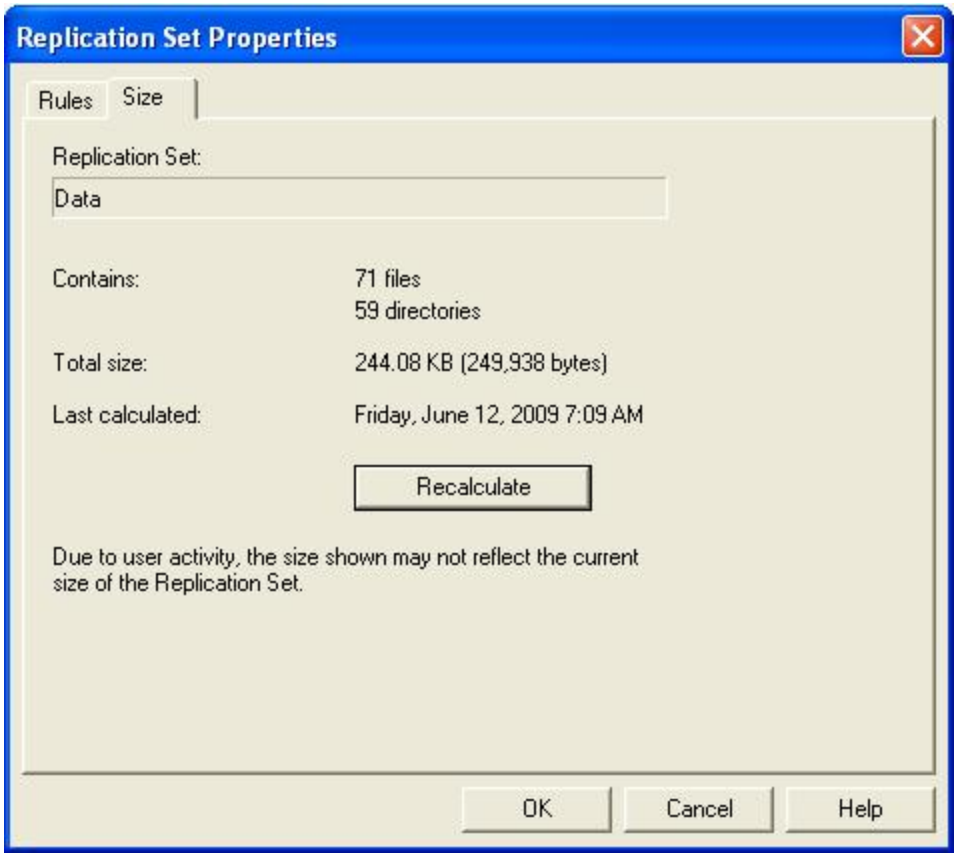


If you save changes to a connected replication set, it is recommended that you perform a mirror to guarantee data integrity between the source and target machines. A dialog box will appear instructing you to disconnect and reconnect the replication set and perform a difference mirror.

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. Right-click on the replication set icon and select **Properties**. The Replication Set Properties dialog box appears.
2. Select the **Size** tab.



3. 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.
4. Click **OK** to return to the Replication Console.



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.

Exporting and importing a replication set

To help reuse replication sets between servers, you can export an existing replication set on one server and import it on another.

- **Exporting a replication set**—Right-click an existing replication set and select **Export**. Select a location and file name for the replication set information, and click **Save**. If you want to share the replication set information with other consoles, select a location accessible by other consoles.
- **Importing a replication set**—Right-click the server where you want to import the replication set and select **New**, **Import Replication Set**. Locate the replication set information file and click **Open**. By default, the new replication set will have the same name as the original replication set. If desired, modify the name. Press Enter to accept the replication set name. By default, the new replication set is imported in an unsaved state. An unsaved replication set icon is red. Modify the replication set definition (include or exclude volumes or files) and then save the replication set by right-clicking on it and selecting **Save**. A saved replication set icon is black.

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, 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, 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 from 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**.



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.

Chapter 12 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.

- 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.

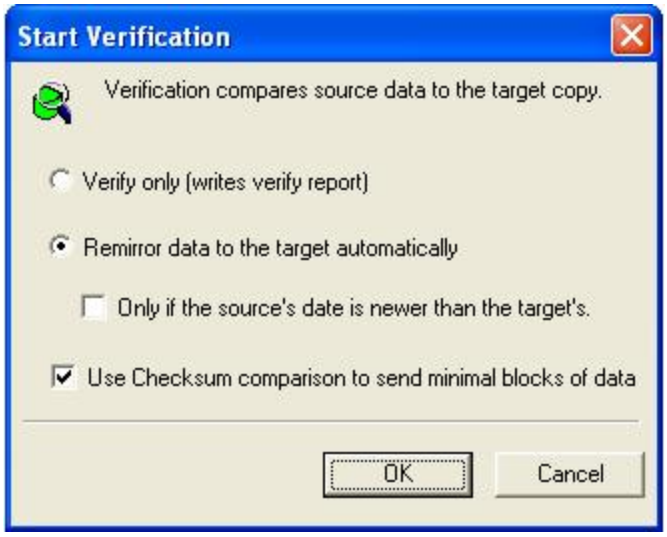


Differences in files on the source and target should be expected for files and applications that are in use during the verification process.

Verifying manually

A manual verification can be run anytime a mirror is not in progress.

1. Right-click the connection on the right pane of the Replication Console and select **Verify**.
2. 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 data 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's date is newer than the target's**—If you are remirroring your files, you can specify that only files that are newer on the source than the target be remirrored.



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 Checksum comparison to send minimal blocks of data**—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.



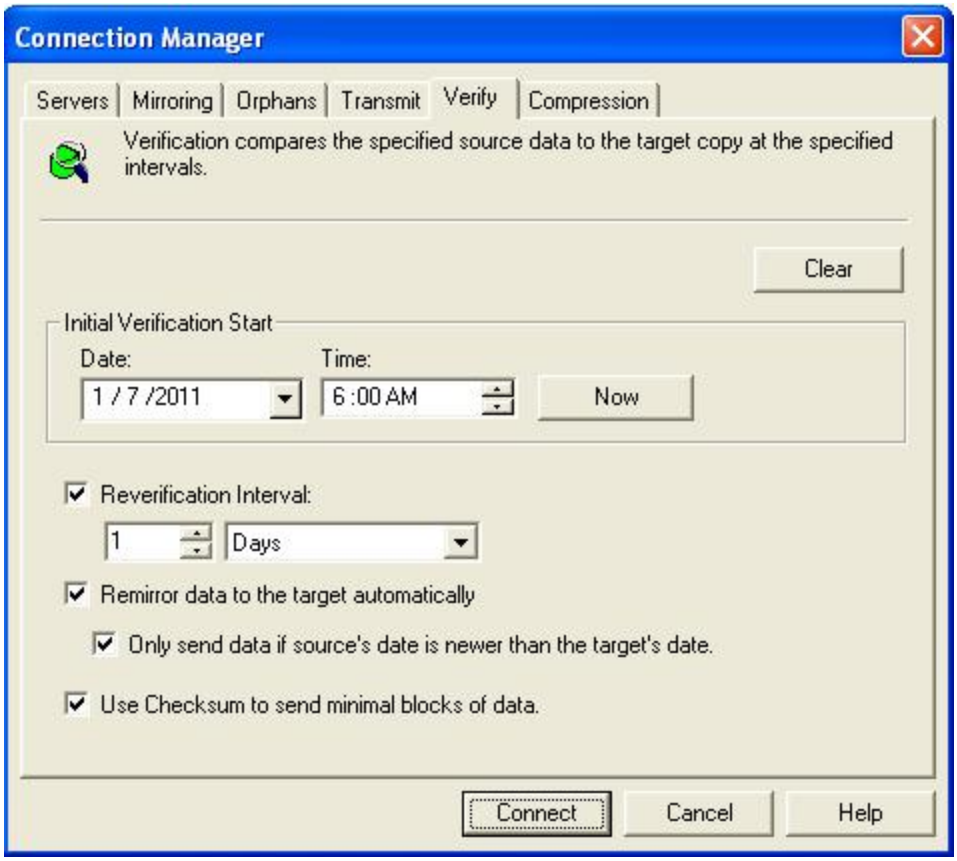
Database applications may update files without changing the date, time, or file size. Therefore, if you are using database applications, you should use the block checksum comparison to ensure proper verification and remirroring.

3. Click **OK** to start the verification.

Verifying on a schedule

Verification can be scheduled to occur automatically at periodic intervals.

1. Right-click the connection on the right pane of the Replication Console and select **Connection Manager**.
2. Select the **Verify** tab.



3. 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.
4. 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.
5. Select if you want to **Remirror data 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.
6. If you are remirroring your files, you can specify **Only send data if source's date**

is newer than the target's date so that only files that are newer on the source than on the target are remirrored.



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.

7. Specify if you want the verification process to **Use Checksum to send minimal blocks of data** 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.



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.

8. 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. Right-click the source server on the left pane of the Replication Console and select **Properties**.
2. Select the **Logging** tab.
3. 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.
4. Under the Verification section, **Filename** contains the base log file name for the verification process. The replication set name will be prepended 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.
5. 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.
6. 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.

   Changes made to the verification log in the **Server Properties**, **Logging** tab will apply to all connections from the current source machine.

7. Specify the **Language** of the log file. Currently, English is the only available language.
8. 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.

Chapter 13 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, right-click an established connection and 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.



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.

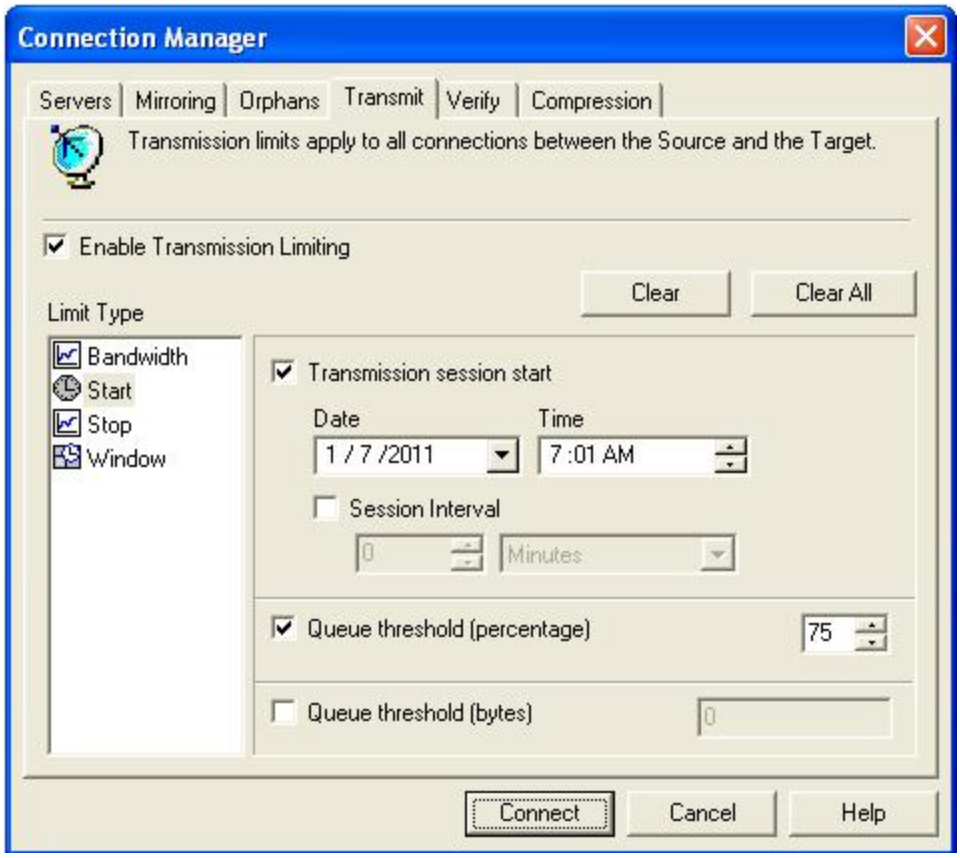
Any replication sets from a source connected to the same IP address on a target will share the same scheduled transmission configuration.

1. Right-click the connection on the right pane of the Replication Console and select **Connection Manager**.
2. 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.
3. 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 p.m., 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

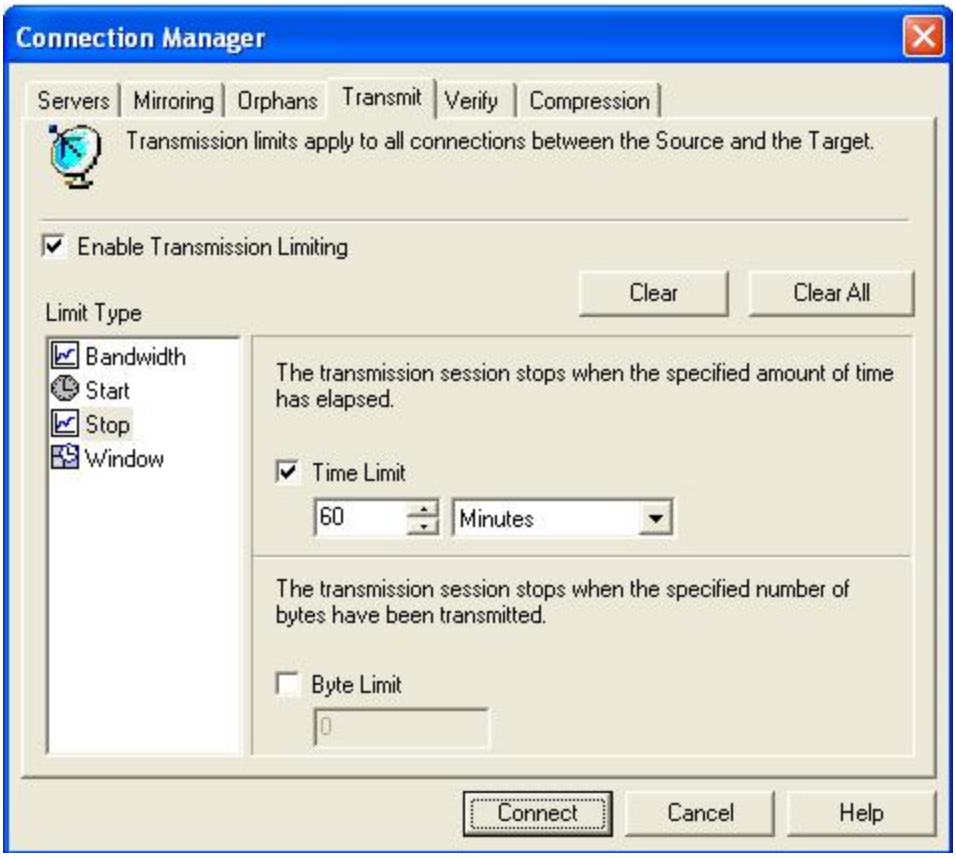
for loss of data. To avoid using the entire queue, you can configure Double-Take Availability to begin transmitting data to the target when the queue reaches a certain point. This point can be defined as a percentage of the disk queue that must be in use or the number of bytes in the disk queue. For 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.



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.

4. 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.

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



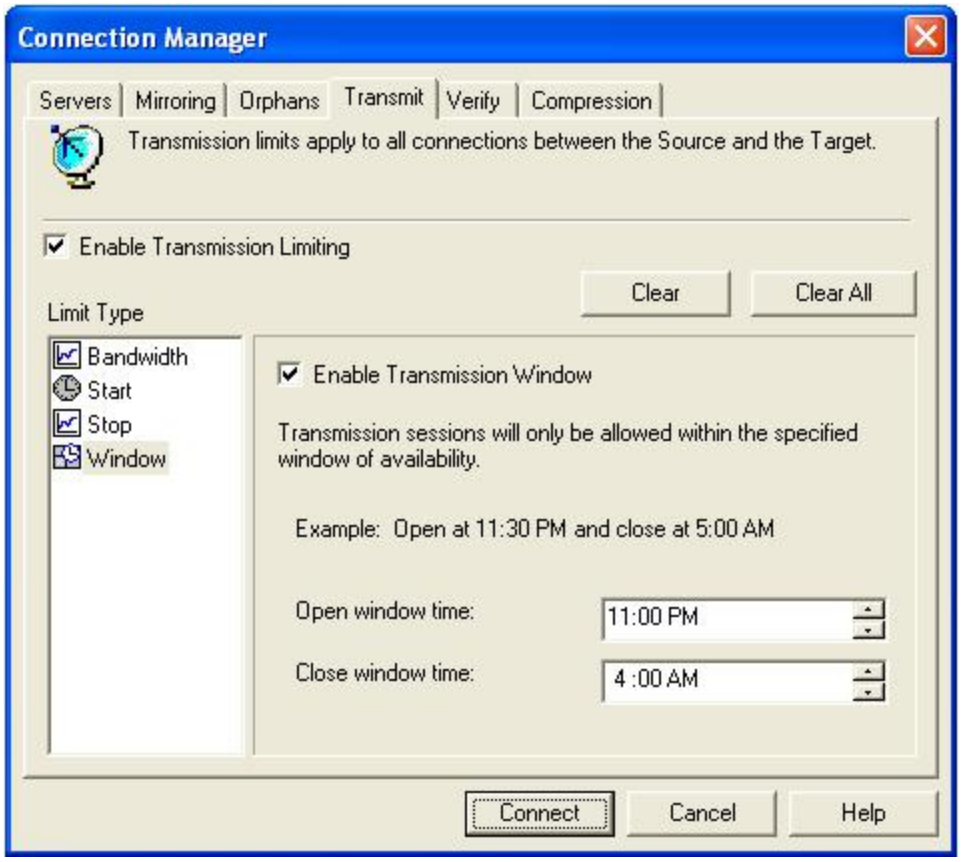
in the queue. The data will be transmitted and when the 10 MB **Byte Limit** is reached, the network connection closes. This is useful in configurations where metered charges are based on connection time.

5. 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.

   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.

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

Limiting transmission bandwidth

Using the Connection Manager **Transmit** tab, you can set start and stop criteria along with a schedule window.



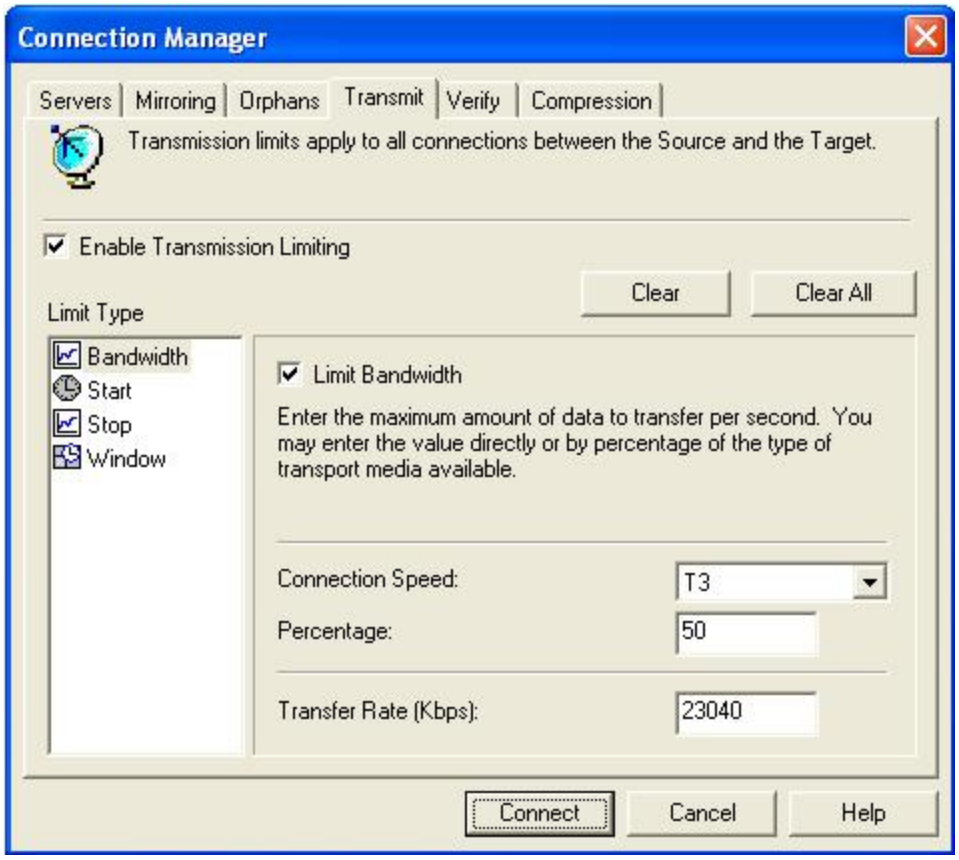
Double-Take Availability checks the schedule once every second, and if a user-defined criteria is met, transmission will start or stop, depending on the option specified.

Any replication sets from a source connected to the same IP address on a target will share the same scheduled transmission configuration.

1. Right-click the connection on the right pane of the Replication Console and select **Connection Manager**.
2. 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.
3. Select the **Bandwidth** option in the **Limit Type** box. Mark the **Limit Bandwidth** check box to enable the bandwidth limiting features. Define the bandwidth available for Double-Take Availability transmission by using either of the following options.



- **Percentage**—Specify the percentage of bandwidth to be used for Double-Take Availability transmissions and the total bandwidth capacity that is available.
- **Transfer Rate**—Specify the number of kilobits to send every second.

The only value that is persistently stored is the number of kilobits per second. When the page is refreshed, the percentage and available bandwidth capacity may not be the same value that you entered. Double-Take Availability changes these values to the maximum values for the smallest possible link.

---

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

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.



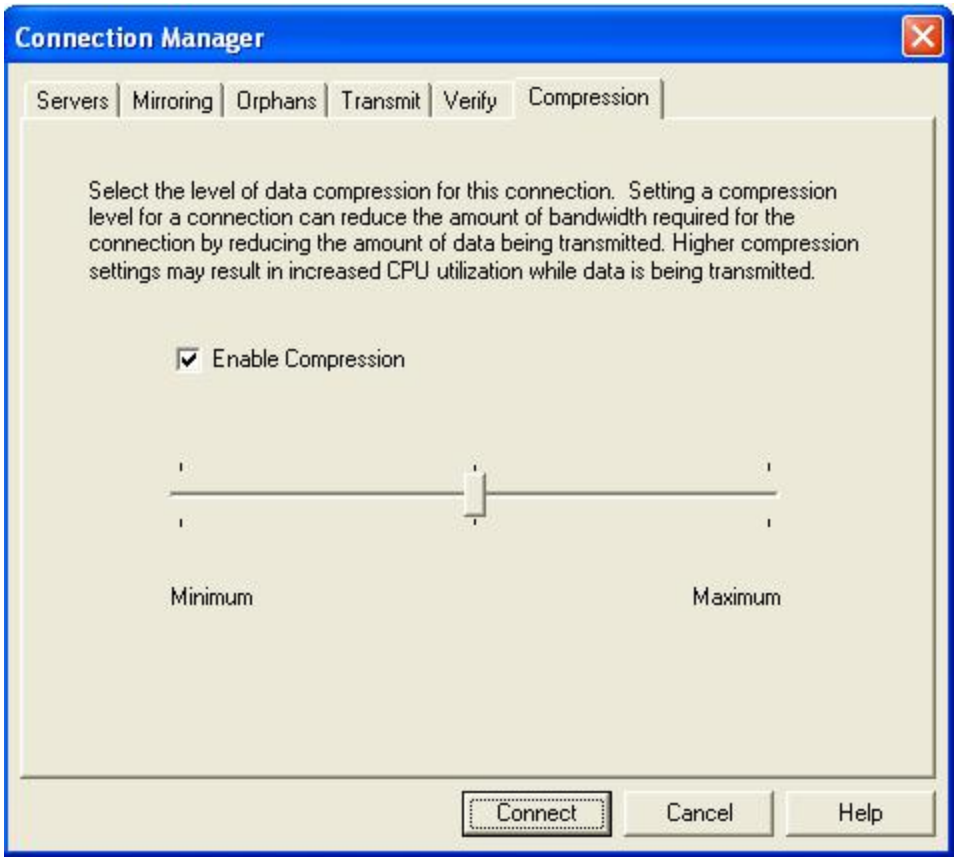
Any replication sets from a source connected to the same IP address on a target will share the same compression configuration.

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 for setting 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.

Chapter 14 Failover and failback

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 machine 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 machine 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 machine that only provides processing services, such as a web server.

Failover can be configured to stand in for one or more IP addresses associated with different NICs on the source. Each IP address can be added to a specific target NIC making NIC configuration very flexible. For example, a single NIC on the source may have one or more IP addresses assigned to it. If that source or the NIC fails, all traffic from the source is directed to the target. If there are multiple NICs on the source, the target can assume the traffic from all of the addresses. Additional NICs on the target increase flexibility and control. Secondary target NICs can assume the traffic from a failed source NIC while normal target traffic can continue to use the primary target NIC.

Failback is the process in which the target releases the source identity so that the source can be brought back onto the network.

- Configuring failover monitoring
- Editing failover monitoring configuration
- Monitoring failover
- Failing over
- Removing failover monitoring configuration

Configuring failover monitoring

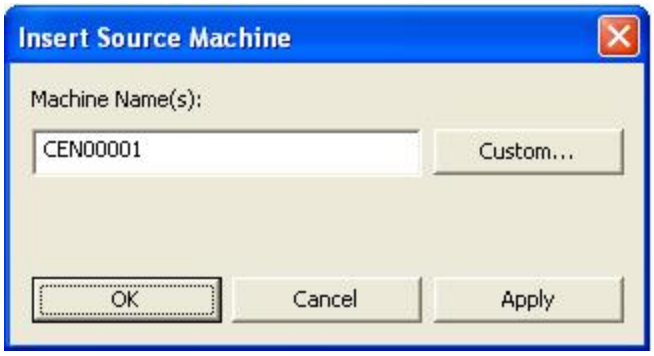
Before beginning your failover configuration, review your IP address and subnet configuration on the source. Because of limitations in the way the Linux kernel handles IP address aliases, you will not be able to mix subnets on the eth0 network interface. Failover should not cause problems in this configuration, but you will lose IP addresses during failback. Therefore, if you must mix subnets on a single interface, use eth1 or higher.

1. The Failover Control Center can be started from within the Replication Console or from the Windows desktop.
   - From the Replication Console, select **Tools**, **Failover Control Center**.
   - From the Windows desktop, select **Start**, **Programs**, **Double-Take for Linux**, **Availability**, **Double-Take Failover Control Center**.
2. Select a failover target from the **Target Machine** list box.

   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 either of the following methods.



   - Type the name of the machine that you want to monitor in **Machine Name(s)** 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.

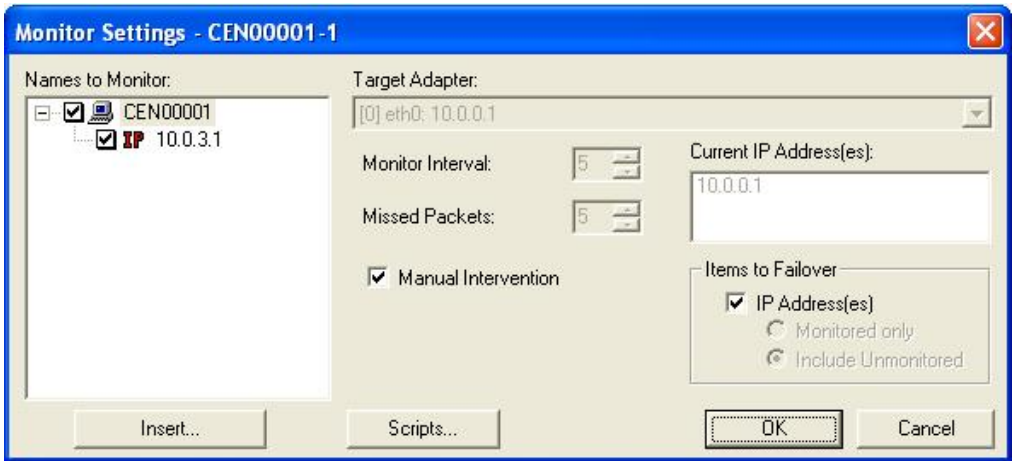


Current IP Addresses displays the IP address(es) currently assigned to the selected target adapter.

8. Highlight an IP address that you have selected for monitoring and select a **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. Repeat this step for each IP address that is being monitored.
9. 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. Repeat this step for each IP address that is being monitored.



To achieve shorter delays before failover, use lower **Monitor Interval** and **Missed Packets** values. This may be necessary for IP addresses on machines, 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 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.



10. 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. When the source and target are on the same subnet, generally a LAN environment, you should failover the IP address. If the source and target are on different subnets, generally a WAN environment, you should not failover the IP address. See WAN considerations for options on handling WAN failover.
        - **Monitored only**—Only the IP address(es) that are selected for monitoring will be failed over.
        - **Include Unmonitored**—All of the IP address(es) will be failed over.

        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.

11. 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.
12. 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 Linux command, executable, or script file. Examples of functions specified in scripts include stopping daemons on the target before failover because they may not be necessary while the target is standing in for the source, stopping daemons on the target that need to be restarted with the source's machine name and IP address, starting daemons 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 daemons on the target after failback because they are no longer needed, stopping daemons 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.
13. If you want to delay the failover or failback processes until the associated script has completed, mark the appropriate check box.
14. If you want the same scripts to be used as the default for future monitor sessions, mark the appropriate check box.
15. Click **OK** to return to the Monitor Settings dialog box.
16. Click **OK** on the Monitor Settings dialog box to save your monitor settings and begin monitoring for a failure.

WAN considerations

When the source and target are on the same subnet, generally a LAN environment, you should failover the IP address. However, if the source and target are on different subnets, generally a WAN environment, you should not failover the IP address. You have several options for handling WAN failover.

- **DNS updates**—You can script DNS updates to modify, at failover time, the source server's DNS A records to have the IP address of the target. When clients resolve a name to an IP address, they will resolve to the target IP address. Depending on the domain size and how DNS updates are propagated, it may take several minutes or even hours for the updates to complete.
- **Reconfigure routers using a failover script**—You can automatically reconfigure routers using a failover script to move the source's subnet from the source's physical network to the target's physical network. 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.

- **VPN infrastructure**—A VPN infrastructure allows your source and target to be on the same subnet, in which case IP address failover will work the same as a LAN configuration.

Scripting example using the BIND DNS client

1. Install the BIND DNS client on the target server, if it is not already installed.
2. Create a PATH statement on the target for the BIND directory to ensure that it runs every time the executable is called.
3. Update the Double-Take daemon on the target to use a domain account that has rights to modify BIND DNS. You will have to stop and restart the daemon for changes to the user account to take effect.
4. Create a failover script with the following command. Specify this script for post-failover.

```
nsupdate "dnsover"
```

5. Create a file called dnsover and add the following lines. This is the file called by your post-failover script.

```
# Substitute your source name, target name, and target IP address

update delete source_server_name.fully_qualified_domain.com
update add target_server_name.fully_qualified_domain.com 86400 A target_server_IP
send
```

6. Create a failback script with the following command. Specify this script for post-failback.

```
nsupdate "dnsback"
```

7. Create a file called dnsback and add the following lines. This is the file called by your post-failback script.

```
# Substitute your target name, source name, and source IP address
update delete target_server_name.fully_qualified_domain.com A
update add source_server_name.fully_qualified_domain.com 86400 A source_server_IP
send
```

When failover and failback occur, the failover and failback scripts will automatically trigger DNS updates.

Protecting NFS shares

A share is any local volume, drive, or directory resource that is exported and shared across a network. NFS shares must be configured for failover through the failover scripts or created manually on the target.

1. Mount DTFS on the directory on the source where the NFS share is located.
2. Start the Double-Take daemon on the source.
3. Stop and restart the NFS daemon on the source. The Double-Take daemon must be running before the NFS daemon in order for replication operations to be captured.
4. On your target, set the NFS daemon to manual startup. This allows the failover script to control when the daemon starts on the target.
5. Create a replication set on the source that includes /etc/exports and the shared data that has DTFS mounted.
6. Connect the replication set using the Connection Wizard or the Connection Manager.
7. Add the following to your post-failover script.

```
service nfs start
```

8. If necessary, update DNS.

After failover, the NFS daemon will automatically be started by the post-failover script. If your clients see a stale file handle error message when attempting to access a share, they will need to reconnect to the share.

Protecting Samba shares

A share is any local volume, drive, or directory resource that is exported and shared across a network. Samba shares must be configured for failover through the failover scripts or created manually on the target.

1. Mount DTFS on the directory on the source where the Samba share is located.
2. Start the Double-Take daemon on the source.
3. Stop and restart the Samba daemon on the source. The Double-Take daemon must be running before the Samba daemon in order for replication operations to be captured.
4. On your target, set the SMB and WinBind daemons to manual startup. This allows the failover script to control when the daemons start on the target.
5. Create a replication set on the source that includes /etc/SAMBA/samba_conf and the shared data that has DTFS mounted.
6. Connect the replication set using the Connection Wizard or the Connection Manager.
7. Add the following to your post-failover script.

```
service smb start
```

8. If necessary, update DNS.

After failover, the daemons will automatically be started by the post-failover script. If your clients see an access denied or share not found error message when attempting to access a share, they will need to remount the share.

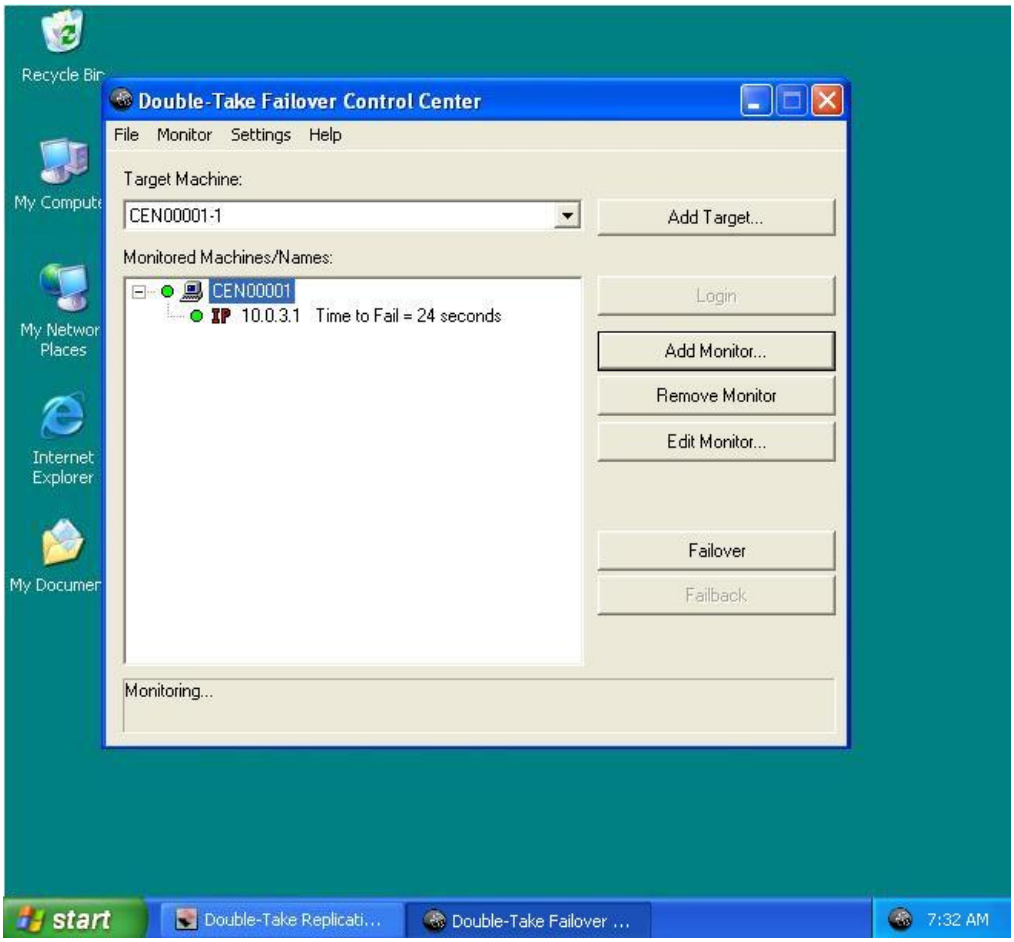
Editing failover monitoring configuration

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 instructions under Configuring failover monitoring.

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



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.

Failing over

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.



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 Linux system 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.

Files that were open or being accessed at the time of failover will generate Stale NFS file handle error messages. Remount the NFS share to correct this error.

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.

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.

Chapter 15 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 other method allows you to failback first and then restore the data from the target to the source. This method may be easier in some situations, but 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 processes.

- Restoring then failing back
- Failing back then restoring

Restoring then failing back

Use these instructions to restore your data first and then failback.

1. Resolve the problem(s) on the source that caused it to fail. If you have to rebuild your source, use a unique identity.
2. Stop any applications that were failed over that may be running on your source. The files must be closed on the source so that updated files from the target will successfully overwrite the files on the source during the restoration.
3. Modify the source so that it can be brought onto the network with a new, unique IP address or one that was not failed over. It needs to be able to exist on the network without an IP address conflict and communicate with the target.
4. 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.
   - The source and target can communicate with each other.
   - All applications on the source are stopped.
5. From your target, confirm the Replication Console is communicating with the source using the new, unique 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 by the new IP address and click **OK**.
6. Disconnect the connection from the original source to the target, if it still exists.
7. Begin your restoration process.
   a. From the Replication Console, select **Tools**, **Restoration Wizard**.
   b. Review the Welcome screen and click **Next** to continue.

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

   c. Select the target that contains the current copy of the data that you want to restore and click **Next**.
   d. Select the original source or **Alternate**, if your original source is not listed. This option identifies to the target which data set you are trying to restore so that the appropriate replication sets can be presented to you.
   e. Click **Next** to continue.

f. Specify if you want to use an existing replication set or create a new one. This replication set will be used to connect from the target to the source.
   - **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. You will have an opportunity to modify the replication set definition.
   - **Create a new replication set with this name**—If you choose to create a new replication set, specify a replication set name.With this option, you will need to define the data to be restored.

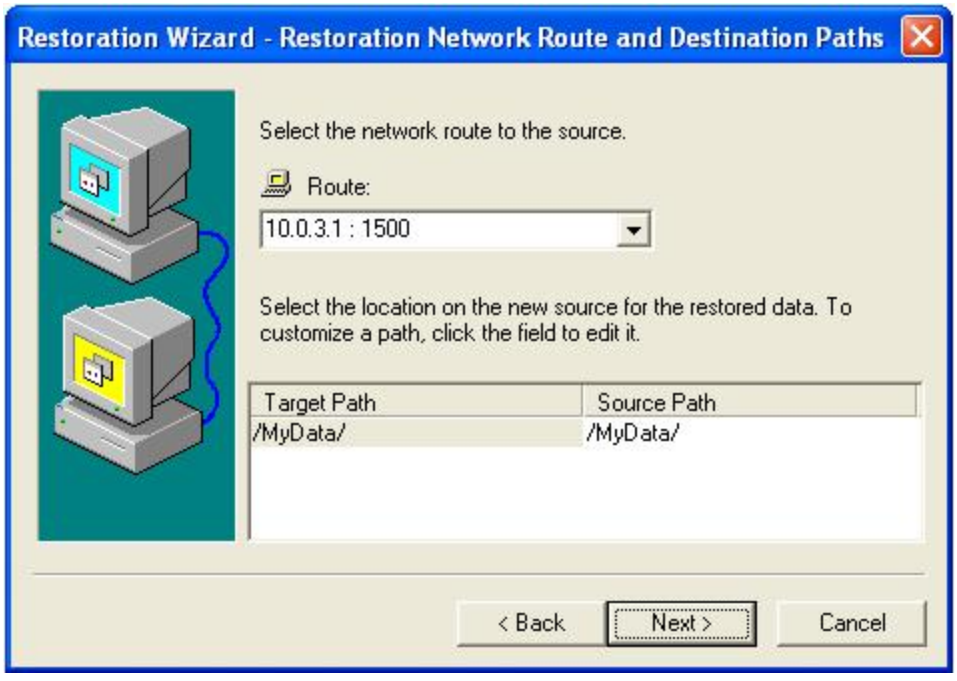
g. Click **Next** to continue.

h. A tree display appears identifying the data available for restoration. Mark the check box of the volumes, directories, and/or files you want to restore. Keep in mind that if you exclude volumes, folders, and/or files that were originally replicated, it may compromise the integrity of your applications or data.

i. Click **Next** to continue

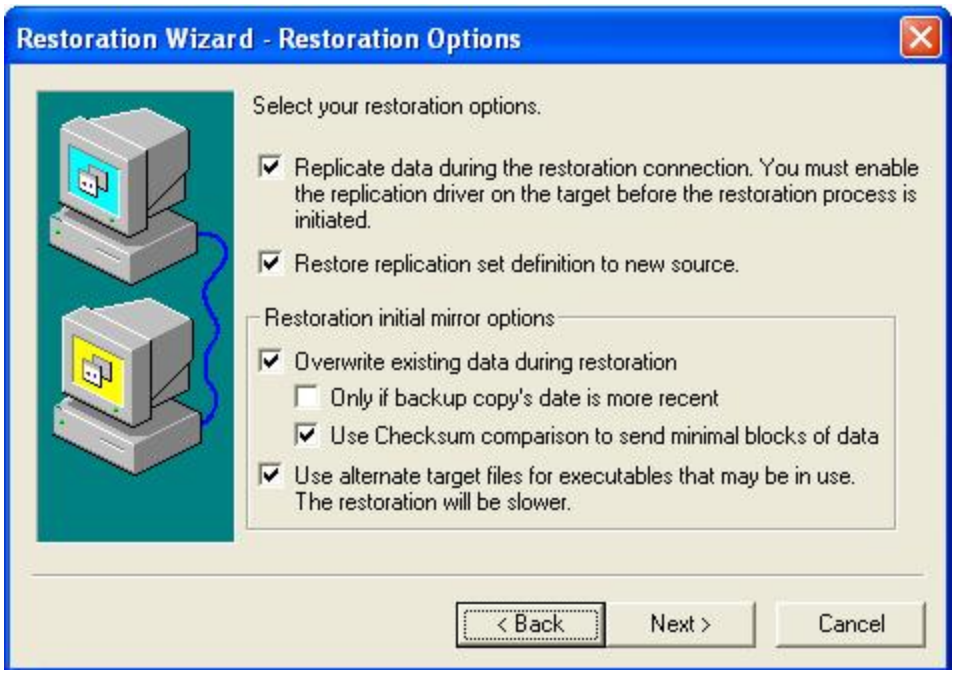
j. Select the new source server. This is the server where the data from the target will be restored. This may be the original source server or a new server. Click **Next** to continue.

k. Select your network route to the new source, which includes the IP address and port number. Also select the location on the new source for the restored data. If you want to set a customized path, click in the field under **Source Path** to edit the location.



l. Click **Next** to continue.

m. Specify the restoration options that you want to use.



- **Replicate data during the restoration connection**—This option allows you to replicate on-going data changes during and after the restoration mirror is performed. Use this option if the source data on the target will continue to change during the restoration process. You do not need to use this option if the source data on the target is not changing. If you do not select this option, any data changes that might occur on the target after the restoration process is initiated will not be transmitted to the source. If you do select this option, you must configure replication on the target prior to initiating the restoration process.
- **Restore replication set definition to new source**—This option restores a copy of the replication set database on the target to the new source.
- **Overwrite existing data during restore**—This option restores all existing files by overwriting them and writes any files that do not exist. If this option is disabled, only files that do not exist on the new source will be restored.
    - **Only if backup copy's date is more recent**—This option restores only those files that are newer on the target than on the new source. The entire file is overwritten with this option.



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 restored.

- **Use Checksum comparison to send minimal blocks of data**—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 new source.



To ensure data integrity, the replicate during restoration and overwrite existing data options are dependent on each other. If you want to enable replication, overwrite data will automatically be enabled. If you disable the option to overwrite data, replication will automatically be disabled.

- **Use alternate target files for executables that may be in use**—If you have executables that may be in use during the restoration, you can have Double-Take Availability create and update an alternate file during the restoration. Once the mirroring and replication operations have been completed, the alternate file will be renamed to the original file. This process will reduce the speed of your restoration, so it should only be used if executables may be in use.

n. Review your selections and click **Finish** to begin the restoration.

8. Monitoring the restoration connection and 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.
9. When you are ready, begin the failback process.
    a. Stop user access to the target.
    b. 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.
    c. Disconnect the restoration connection.
    d. Open the Failover Control Center.
    e. Select the original target that is currently standing in for the original failed source.
    f. Highlight the failed source and click **Failback**. The user downtime starts now. If you have a pre-failback script configured, it will be started.

g. 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 failover monitoring, do not select either option. Leave the prompt dialog box open as is.

10. On the source, modify the identity back to the original source IP address.
11. 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 by the original IP address and click **OK**.
12. 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.



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.

At this time, you can start any applications and allow end-users to access the data.

Failing back then restoring

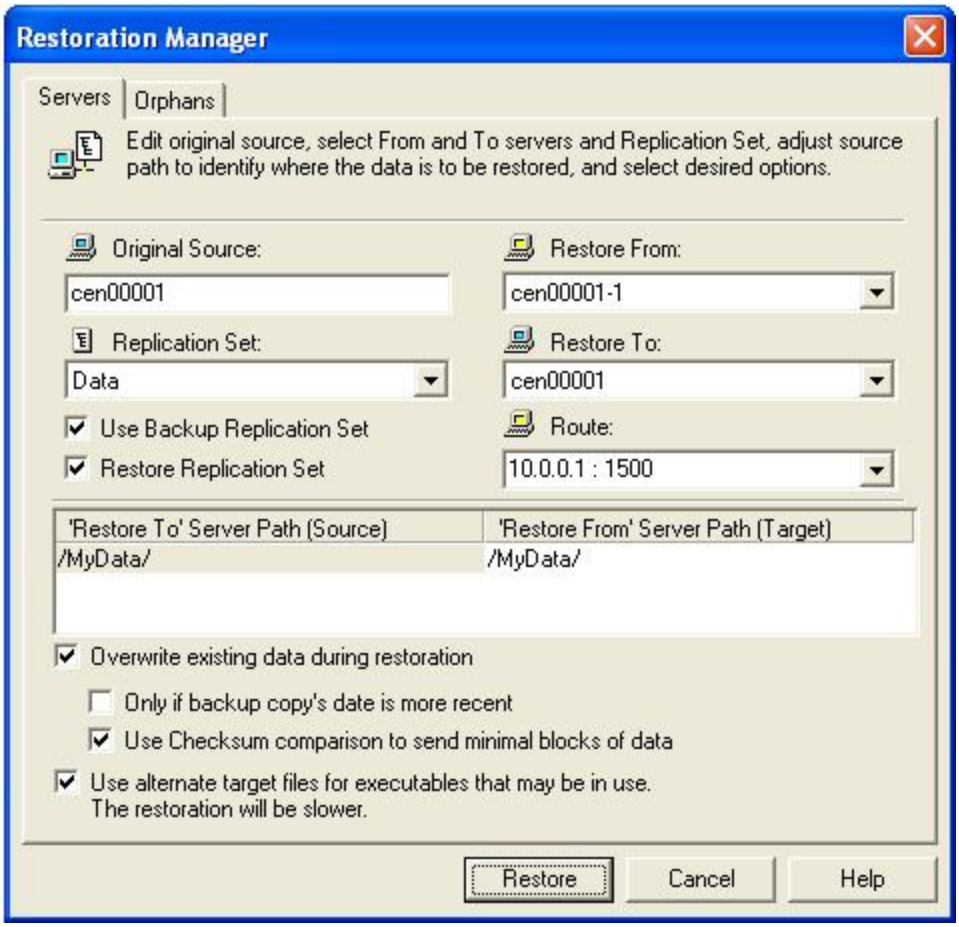
Use these instructions to failback first and then restore your data.

1. Resolve the problem(s) on the source that caused it to fail. If you have to rebuild your source, make sure you use the same identity as the original source configuration.
2. Because you do not want your users accessing the source or its data until newer data from the target can be restored, deny access to user logins by setting /etc/nologin. See your Linux documentation for details on creating this file.
3. 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 during the restoration.
4. From the Failover Control Center, select the target that is currently standing in for the failed source.
5. Select the failed source and click **Failback**. The user downtime starts now. If you have a pre-failback script configured, it will be started.
6. 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. 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 machine, the source post-failback script, if configured, will be started.

   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. From the Replication Console, select **Tools**, **Restoration Manager**.



8. Identify the **Original Source** machine. This is your source machine where the data originally resided.
9. Select the **Restore From** machine. This is the target machine where the copy of the data is stored.
10. **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.
11. Select the **Restore To** machine. This is your source where the updated data from the target will be sent.
12. 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.
13. 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.
14. Select the **Route** on the target. This is 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.

15. The **Restore To Server Path** and **Restore From Server Path** 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.

    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.

16. Select the restoration conditionals that you want to use.
    - **Overwrite existing data 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's date 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.

      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 Checksum comparison to send minimal blocks of data**—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.
17. 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.
18. Click **Restore** to begin the restoration.

After the restoration is complete, the restoration connection will automatically be disconnected and the replication set deleted. At this time, you can start any applications and allow end-users to access the data on the source.

Chapter 16 Restoration

If you are not using failover, you have the choice of restoring specific volumes, directories, and/or files, or you can restore an entire server. If you are restoring an entire server, you can choose to restore it to similar hardware or to an ESX virtual server.

- Restoring files only
- Restoring an entire server
- Restoring an entire server to an ESX virtual server

Restoring files only

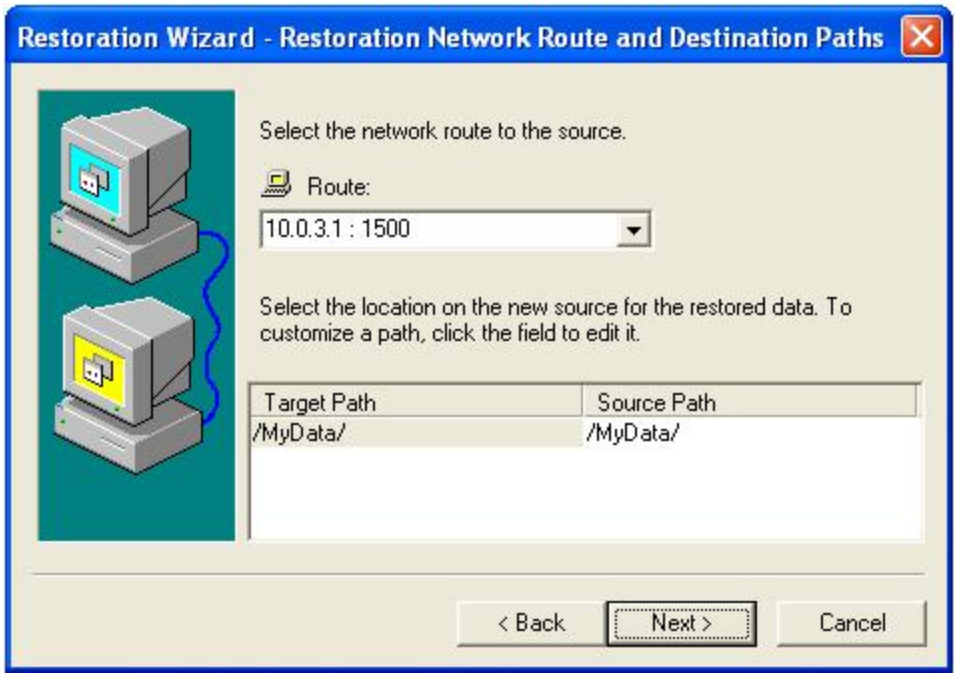
Use these instructions to restore specific volumes, directories, and/or files.

1. From the Replication Console, select **Tools**, **Restoration Wizard**.
2. Review the Welcome screen and click **Next** to continue.

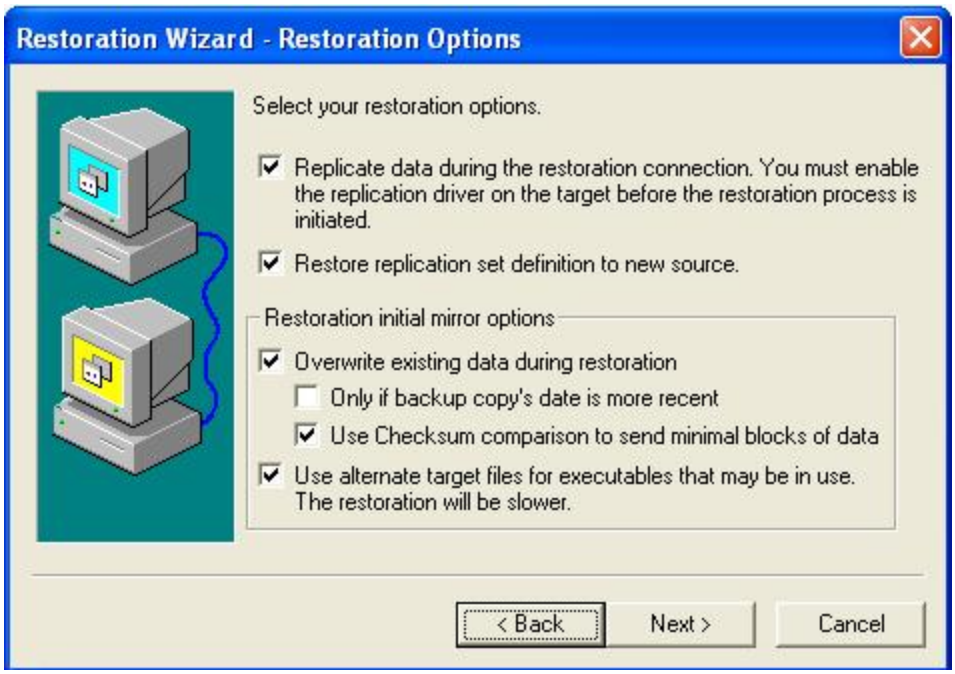
   

   At any time while using the Restoration Wizard, click **Back** to return to previous screens and review your selections.

3. Select the target that contains the copy of the data that you want to restore and click **Next**.
4. Select the original source or **Alternate**, if your original source is not listed. This option identifies to the target which data set you are trying to restore so that the appropriate replication sets can be presented to you.
5. Click **Next** to continue.
6. Specify if you want to use an existing replication set or create a new one. This replication set will be used to connect from the target to the source.
   - **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.
   - **Create a new replication set with this name**—If you choose to create a new replication set, specify a replication set name.
7. Click **Next** to continue.
8. A tree display appears identifying the data available for restoration. Mark the check box of the volumes, directories, and/or files you want to restore. Keep in mind that if you exclude volumes, folders, and/or files, it may compromise the integrity of your applications or data.
9. Click **Next** to continue
10. Select the new source server. This is the server where the data from the target will be restored. This may be the original source server or a new server. Click **Next** to continue.
11. Select your network route to the new source, which includes the IP address and port number. Also select the location on the new source for the restored data. If you want to set a customized path, click in the field under **Source Path** to edit the location.



12. Click **Next** to continue.
13. Specify the restoration options that you want to use.



- **Replicate data during the restoration connection**—This option allows you to replicate on-going data changes during and after the restoration mirror is performed. Use this option if the source data on the target will continue to change during the restoration process. You do not need to use

this option if the source data on the target is not changing. If you do not select this option, any data changes that might occur on the target after the restoration process is initiated will not be transmitted to the source. If you do select this option, you must configure replication on the target prior to initiating the restoration process.

- **Restore replication set definition to new source**—This option restores a copy of the replication set database on the target to the new source.
- **Overwrite existing data during restore**—This option restores all existing files by overwriting them and writes any files that do not exist. If this option is disabled, only files that do not exist on the new source will be restored.
  - **Only if backup copy's date is more recent**—This option restores only those files that are newer on the target than on the new source. The entire file is overwritten with this option.

    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 restored.

  - **Use Checksum comparison to send minimal blocks of data**—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 new source.

    To ensure data integrity, the replicate during restoration and overwrite existing data options are dependent on each other. If you want to enable replication, overwrite data will automatically be enabled. If you disable the option to overwrite data, replication will automatically be disabled.

- **Use alternate target files for executables that may be in use**—If you have executables that may be in use during the restoration, you can have Double-Take Availability create and update an alternate file during the restoration. Once the mirroring and replication operations have been completed, the alternate file will be renamed to the original file. This process will reduce the speed of your restoration, so it should only be used if executables may be in use.

14. Review your selections and click **Finish** to begin the restoration.

After the restoration connection is established, you can right-click on it to access the Connection Manager and advanced connection options. If you selected to **Replicate**

data during the restoration connection, the connection will remain in place, replicating any data changes, until you manually disconnect it. When you disconnect it, the replication set created for the restoration connection will automatically be deleted. If you did not select to replicate data, the restoration connection will automatically be disconnected and the replication set deleted when the restoration mirror is complete.

Restoring an entire server

Use these instructions to restore an entire server to similar hardware.

1. For Red Hat-based servers, reboot the server you are restoring to into single user mode. This may be the original source server or a new server. For Novell-based servers, you will want to reboot to run-level 2. You cannot issue a run-level change command. You must reboot in order to pause necessary processes.
    a. At the grub boot menu, press any key.
    b. Highlight your kernel from the list.
    c. For Red Hat-based servers, press **a** to append additional options to the command line, and then at the grub append command prompt, add the word **single** to the end of the command and press Enter.



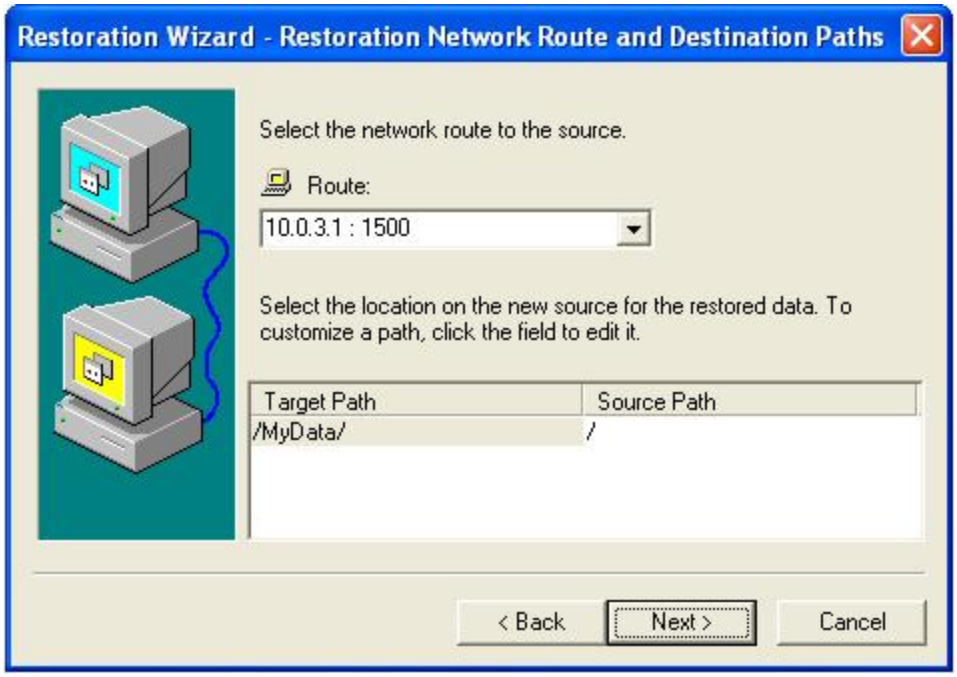
    d. For Novell-based servers, press 2 to take the server to run-level 2 and then manually shutdown non-essential services.
    e. After the boot process has completed, run the following initialization scripts. For Novell-based servers, you do not need to run the network script because run-level 2 includes networking.
        - /etc/init.d/network start
        - /etc/init.d/DT start
        - /etc/init.d/sshd start
2. From the Replication Console, select **Tools**, **Restoration Wizard**.
3. Review the Welcome screen and click **Next** to continue.



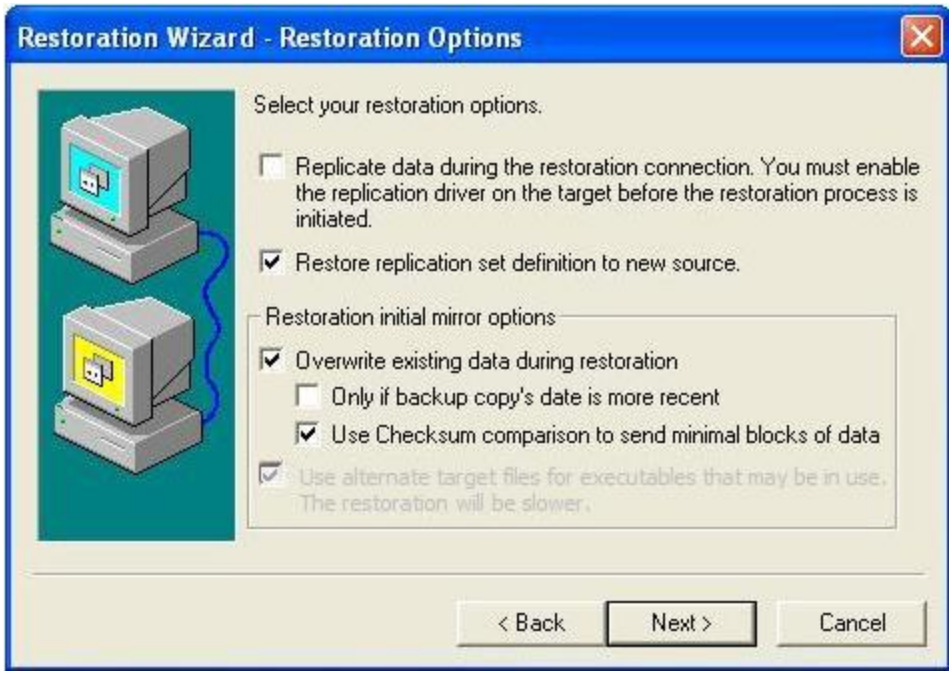
At any time while using the Restoration Wizard, click **Back** to return to previous screens and review your selections.

4. Select the target that contains the copy of the data that you want to restore and click **Next**.
5. Select the original source. This option identifies to the target which data set you are trying to restore so that the appropriate replication sets can be presented to you.
6. Click **Next** to continue.
7. Select **Use this replication set** and specify the **Server_Image** replication set.

8. Click **Next** to continue.
9. A tree display appears identifying the data available for restoration. Verify the existing selection of the volumes, directories, and/or files you want to restore. Keep in mind that if you exclude volumes, folders, and/or files, it may compromise the integrity of your applications or data.
10. Click **Next** to continue
11. Select the new source server. This is the server where the data from the target will be restored. This may be the original source server or a new server. Click **Next** to continue.
12. Select your network route to the new source, which includes the IP address and port number. Do not modify the **Source Path** or **Target Path**.



13. Click **Next** to continue.
14. Specify the restoration options that you want to use.



- **Replicate data during the restoration connection**—Disable this option. It is not required when restoring an entire server.
- **Restore replication set definition to new source**—Enable this option. It restores a copy of the replication set database on the target to the new source.
- **Overwrite existing data during restore**—Enable this option. It restores all existing files by overwriting them and writes any files that do not exist. If this option is disabled, only files that do not exist on the new source will be restored.
    - **Only if backup copy's date is more recent**—Disable this option. It restores only those files that are newer on the target than on the new source. The entire file is overwritten with this option.
    - **Use Checksum comparison to send minimal blocks of data**—You can enable or disable this option as desired. It specifies 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 new source.
- **Use alternate target files for executables that may be in use**—This option is enabled by default and cannot be changed in this context.

15. Review your selections and click **Finish** to begin the restoration.

After the restoration mirror is complete, the restoration connection will automatically be disconnected and the replication set will be deleted. After the restoration replication set is deleted, exit the operating system single user mode by rebooting the server. After the

server restarts in full user mode, you can allow end-users to access the data on the source.

Restoring an entire server to an ESX virtual server

Use these instructions to restore an entire server to an ESX virtual server.

1. Create a virtual guest on your ESX server.
   - Configure the memory and CPUs as needed.
   - For the hard drives, use virtual SCSI disks and the LSI Logic SCSI controller
   - Configure your virtual disks using the same configuration (size and file system type) as the original source you are restoring
2. Install the operating system on the guest, making sure to match the operating system version, kernel version, kernel type, and file system as the original source you are restoring.
3. Install Double-Take Availability on the guest, using a valid Double-Take Availability activation code. You can reuse the activation code that was used on your original source, while the original source is not active.
4. Configure networking on the guest, making sure that it can communicate with the Double-Take Availability target which has the copy of the original source you are restoring.
5. For Red Hat-based guests, reboot the guest into single user mode. For Novell-based guests, you will want to reboot to run-level 2. You cannot issue a run-level change command. You must reboot in order to pause necessary processes.
   a. At the grub boot menu, press any key.
   b. Highlight your kernel from the list.
   c. For Red Hat-based guests, press **a** to append additional options to the command line, and then at the grub append command prompt, add the word **single** to the end of the command and press Enter.



   d. For Novell-based guests, press 2 to take the guest to run-level 2 and then manually shut down non-essential services, such as third-party applications or daemons.
   e. After the boot process has completed, run the following initialization scripts. For Novell-based guests, you do not need to run the network script because run-level 2 includes networking.
      - /etc/init.d/network start

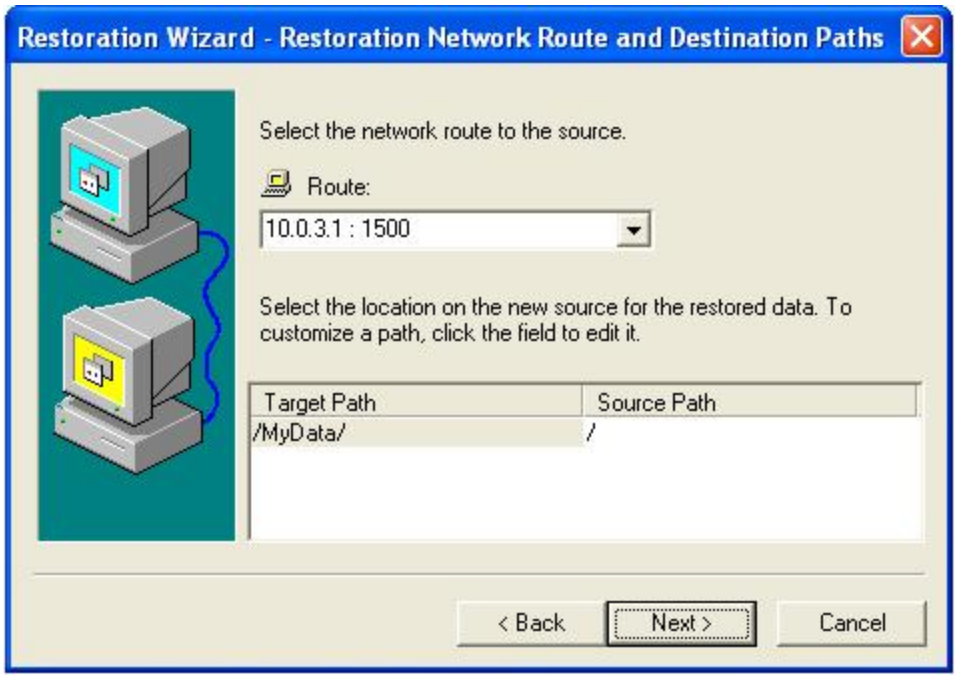
- /etc/init.d/DT start
- /etc/init.d/sshd start

6. Make a copy of the guest's kernel mode configuration file. By default, this is /etc/modprobe.conf for Red Hat-based guests and /etc/sysconfig/kernel for Novell-based guests. Append .pre-migration to the end of the filename when performing the copy. For example, /etc/modprobe.conf will become /etc/modprobe.conf.pre-migration. This copy is required to maintain data specific to your guest that is at risk of being lost by overwriting configuration files during the restoration.
7. From the Replication Console, select **Tools**, **Restoration Wizard**.
8. Review the Welcome screen and click **Next** to continue.

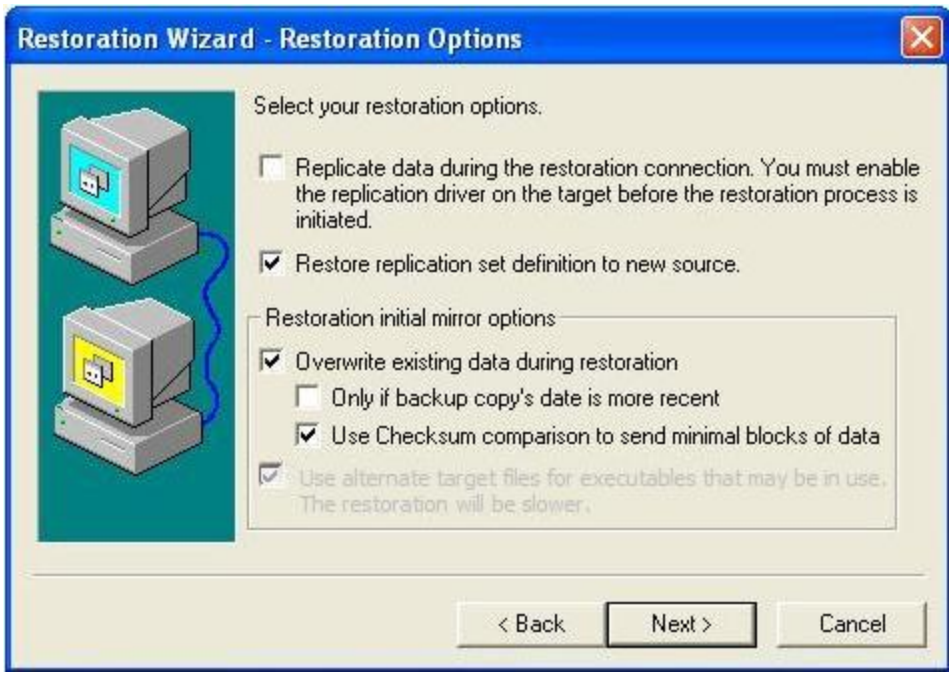


At any time while using the Restoration Wizard, click **Back** to return to previous screens and review your selections.

9. Select the target that contains the copy of the original source that you want to restore and click **Next**.
10. Select the original source. This option identifies to the target which data set you are trying to restore so that the appropriate replication sets can be presented to you.
11. Click **Next** to continue.
12. Select **Use this replication set** and specify the **Server_Image** replication set.
13. Click **Next** to continue.
14. A tree display appears identifying the data available for restoration. Verify the existing selection of the volumes, directories, and/or files you want to restore. Keep in mind that if you exclude volumes, folders, and/or files, it may compromise the integrity of your applications or data.
15. Click **Next** to continue
16. Select the new source server. This is the guest that you just created. Click **Next** to continue.
17. Select your network route to the guest, which includes the IP address and port number. Do not modify the **Source Path** or **Target Path**.



18. Click **Next** to continue.
19. On the Restoration Options page, select the following options.



- **Replicate data during the restoration connection**—Disable this option. It is not required when restoring an entire server.

- **Restore replication set definition to new source**—Enable this option. It restores a copy of the replication set database on the target to the new source.
- **Overwrite existing data during restore**—Enable this option. It restores all existing files by overwriting them and writes any files that do not exist. If this option is disabled, only files that do not exist on the new source will be restored.
  - **Only if backup copy's date is more recent**—Disable this option. It restores only those files that are newer on the target than on the new source. The entire file is overwritten with this option.
  - **Use Checksum comparison to send minimal blocks of data**—You can enable or disable this option as desired. It specifies 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 new source.
- **Use alternate target files for executables that may be in use**—This option is enabled by default and cannot be changed in this context.

20. Review your selections and click **Finish** to begin the restoration.

After the restoration mirror is complete, the restoration connection will automatically be disconnected and the replication set will be deleted. After the restoration replication set is deleted, you will need to perform additional post-restoration steps specific to your guest. If you restored to a Red Hat-based guest, see Configuring Red Hat-based guests after restoring. If you restored to a Novell-based guest, see Configuring Novell-based guests after restoring.

Configuring Red Hat-based guests after restoring

1. SSH to the virtual guest.
2. Confirm that the network scripts are not configured to rely on a hard-coded MAC address. If they are, manually change the scripts to match your virtual server's NIC.
3. Review /boot/grub/device.map to confirm that the configuration matches your virtual hardware. Pay special attention if you are restoring from an IDE or PATA-based host to a SCSI-based virtual server, because the naming conventions between the hardware types are different. For example, if you are restoring from a PATA-based physical server to a virtual server, you will need to change the mapping to read sda instead of hda.
4. Review /etc/fstab and make any necessary adjustments to the file system table to ensure that the file accurately reflects the proper block device location of each file system. Again, be sure to account for differences between PATA and SCSI, as appropriate.
5. Comment out all unnecessary hardware kernel modules that were imported from

your physical host by editing /etc/modprobe.conf. At a minimum, the pcnet32 and mpt* entries must be added to ensure system functionality post-recovery. Once this is complete, run the following command to merge the virtual and physical copies of the file.

cat /etc/modprobe.conf.pre-migration >> /etc/modprobe.conf

6. Run the mkinitrd command to rebuild the kernel's ramdisk to support the changes in hardware. For example, if your physical host used the 2.6.18-128.El kernel, the command would be mkinitrd /boot/initrd.2.6.18-128.EL.img 2.6.18-128.EL.

   mkinitrd <image file> <kernel directory>

7. Reinstall the grub bootloader using the command grub-install --recheck <boot device id>. The boot device ID is typically the device that contains the file system containing /boot, minus the partition ID. For example, if /boot is mounted as /dev/sdb2, the boot device ID is /dev/sdb

8. Reboot the virtual server into single user mode and watch for any file system or other error messages. Kudzu may or may not launch. If it does not, launch it manually to configure any new hardware.



Kudzu is not supplied with Red Hat Enterprise Linux 6.0

Kudzu will prompt for the removal of any hardware devices that no longer exist in the virtual server. Do not remove these items if you plan to return the server to its original hardware.

9. Reboot the server into full user mode.

Configuring Novell-based guests after restoring

1. SSH to the virtual guest.
2. Confirm that the network scripts are not configured to rely on a hard-coded MAC address. If they are, manually change the scripts to match your virtual server's NIC.
3. Review /boot/grub/device.map to confirm that the configuration matches your virtual hardware. Pay special attention if you are restoring from an IDE or PATA-based host to a SCSI-based virtual server, because the naming conventions between the hardware types are different. For example, if you are restoring from a PATA-based physical server to a virtual server, you will need to change the mapping to read sda instead of hda.
4. Review /etc/fstab and make any necessary adjustments to the file system table to ensure that the file accurately reflects the proper block device location of each file system. Again, be sure to account for differences between PATA and SCSI, as appropriate.

5. Compare the INITRD_MODULES line from your physical and virtual kernel configuration files (/etc/sysconfig/kernel and /etc/sysconfig/kernel.pre-migration), adding any declared modules in your pre-restoration file to the current. At a minimum, the pcnet32 and mpt* entries must be added to ensure system functionality post-restoration.
6. Run the mkinitrd command to rebuild the kernel's ramdisk to support the hardware changes.
7. Reconfigure and reinstall the grub bootloader via yast2. Use the yast2 bootloader command, and inspect the properties of each boot menu option to ensure that any drive declarations are up-to-date with your new virtual hardware, and that the correct and matching boot image and kernels are defined.
8. Reboot the virtual server into single user mode, by adding **single** to the boot options line, and watch for any file system or other error messages. If the server is running X Windows, you will need to install the VMware Tools package and configure them for use with X Windows. See your VMware documentation for installation information.
9. Reboot the server into full user mode.

Chapter 17 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 Replication 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*.

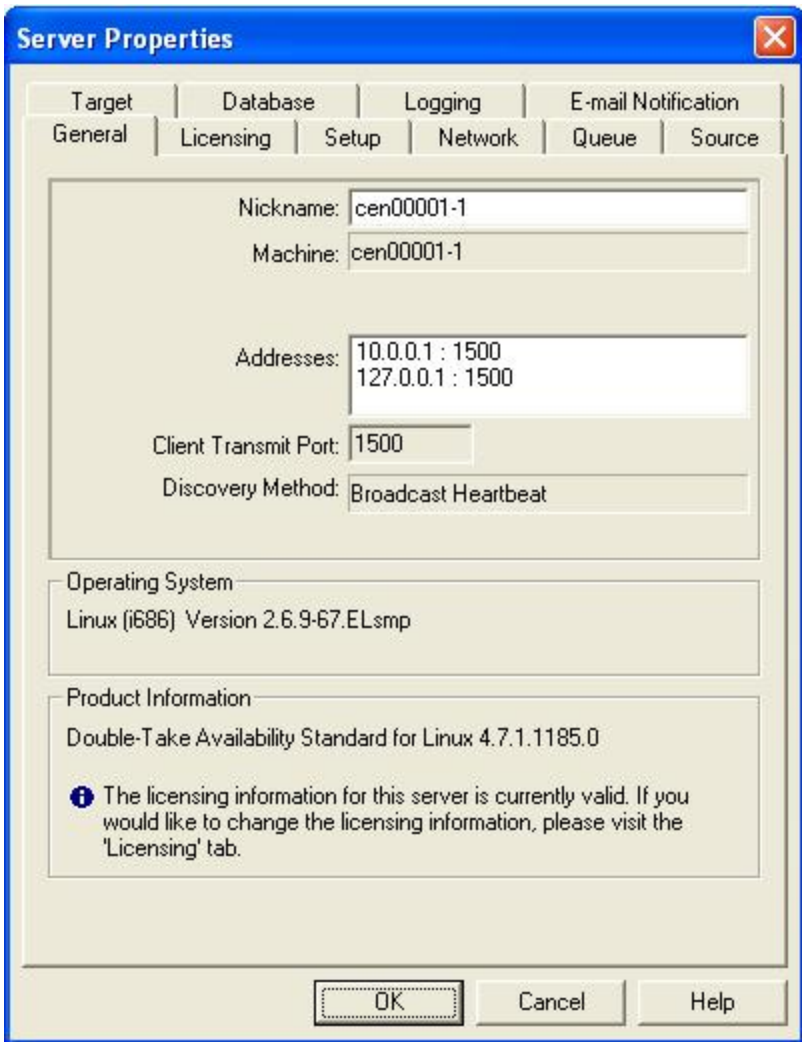
This section contains the following topics, each corresponding to a tab in the Server Properties dialog box.

- 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
- E-mailing system messages

Identifying a server

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

1. Right-click a 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.

- **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**—A Double-Take Availability server was manually inserted into the Replication Console server tree.
  - **Broadcast Heartbeat**—A Double-Take Availability server is broadcasting Double-Take Availability heartbeats.
- **Operating System**—The server's operating system version is displayed.
- **Double-Take Version Information**—The Double-Take Availability version number and build number are displayed.

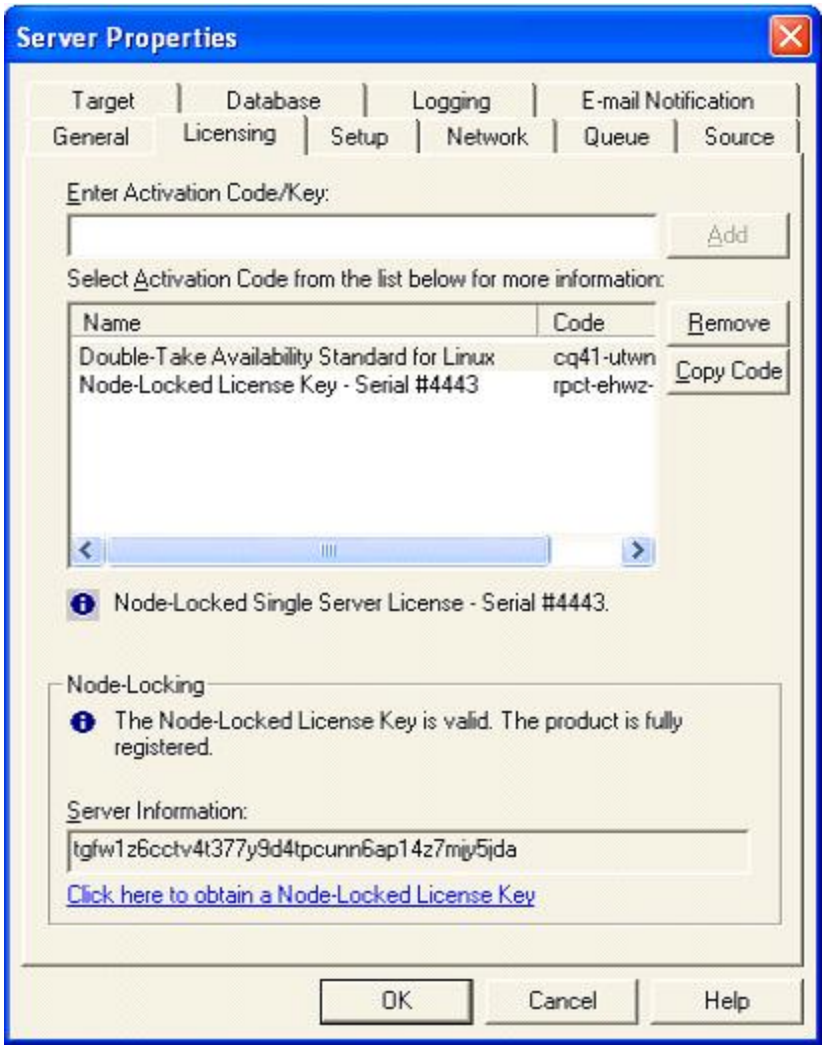
5. Click **OK** to save the settings.

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 have activated it from the Replication Console. 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 Vision Solutions. 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.

1. Right-click a 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.
    a. After entering your temporary node-locked license, click **OK** to activate it. At this point, the temporary license is activated, and you have 14 days to update it to a permanent, node-locked license.
    b. Reopen the Server Properties **Licensing** tab.
    c. Highlight your temporary node-locked license in the list to display the Node-Locking section at the bottom of the **Licensing** tab.

d. Click the hyperlink in the Node-Locking section. If you do not have an Internet connection, copy the **Server Information** text from the Node-Locking section into the form at https://activate.doubletake.com from another machine.

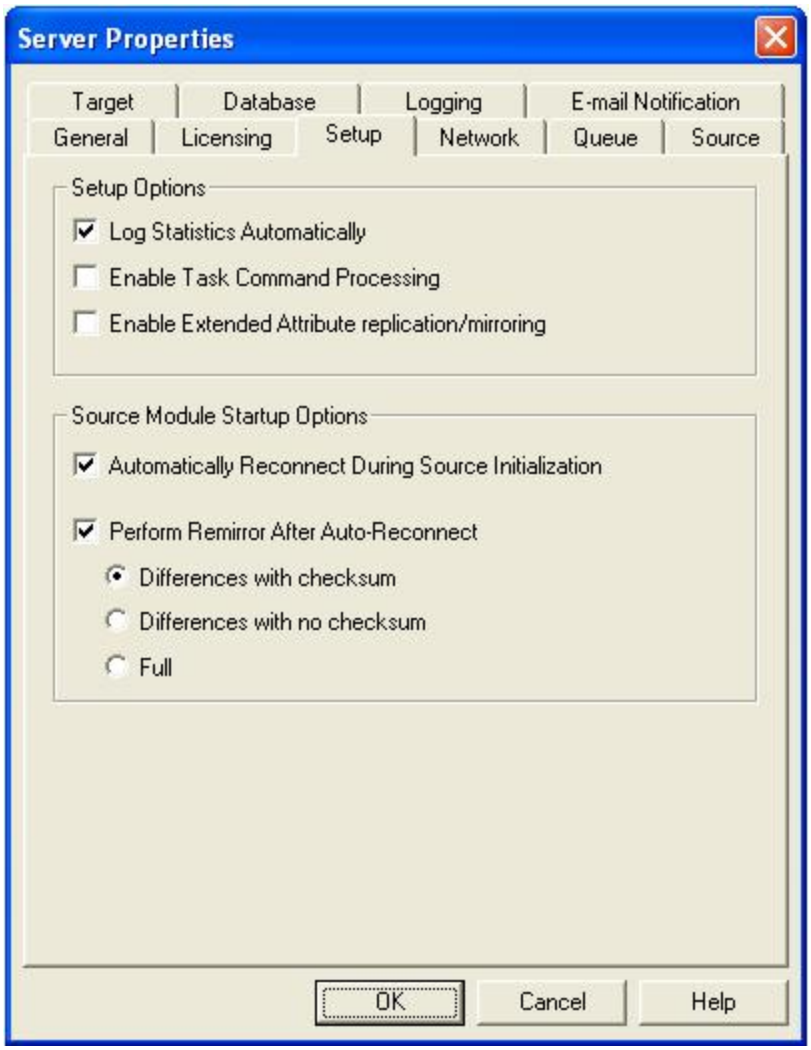
e. After you submit the form, you will receive an email with a node-locked license key. Enter that key on the **Licensing** tab and click **Add**. The permanent activation code is specific to this server. It cannot be used on any other server. If the activation code and server do not match, Double-Take Availability will not run.

8. Click **OK** to apply the activation code(s) you entered and/or the node-locked license key.

Configuring server startup options

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

1. Right-click a 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.

- **Enable Extended Attribute replication/mirroring**—This option enables extended attribute replication and mirroring. You must enable this option on all of your source and target servers.
- **Automatically Reconnect During Source Initialization**—If enabled, Double-Take Availability will automatically reconnect any connections that it automatically disconnected.
- **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.
  - **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.
  - **Full**—All files are sent to the target.

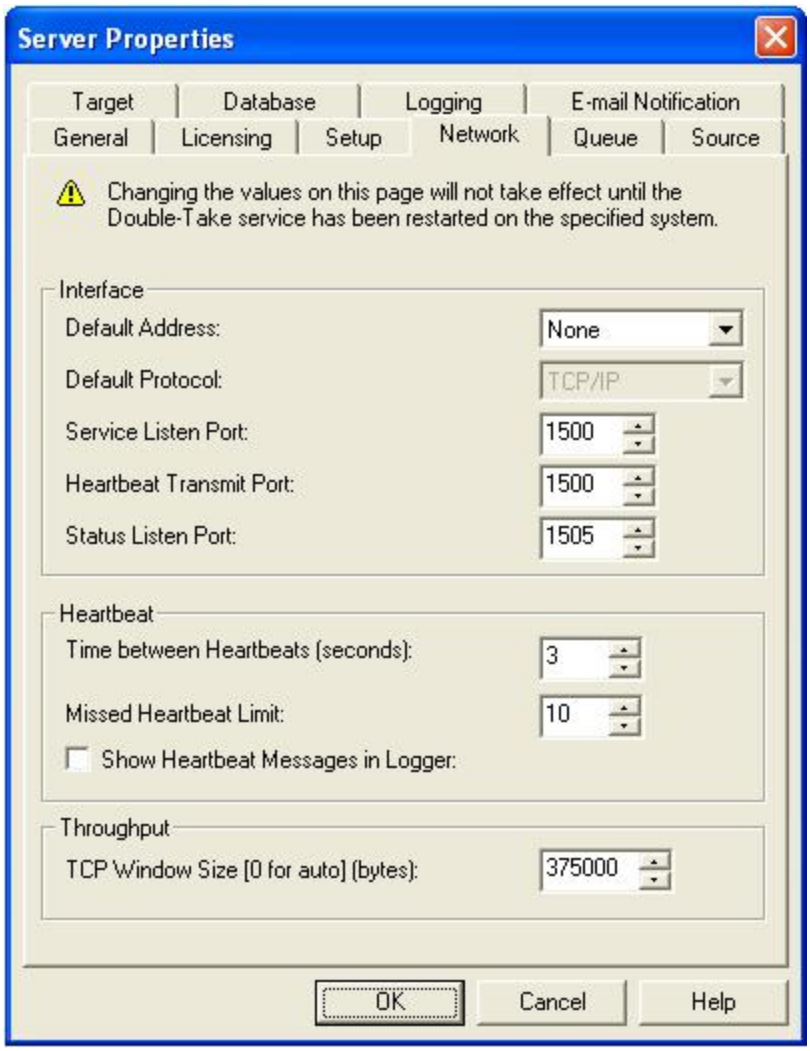


Database applications may update files without changing the date, time, or file size. Therefore, if you are using database applications, you should use the **Differences with checksum** or **Full** option.

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

Configuring network communication properties for a server

1. Right-click a 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.
   - **Default Protocol**—The default protocol for all Double-Take Availability communications is the TCP/IP protocol. In the future, Double-Take Availability may support other communication protocols.

- **Service Listen Port**—Double-Take Availability servers use the **Service Listen Port** to send and receive commands and operations between two Double-Take Availability servers.
- **Heartbeat Transmit Port**—A Double-Take Availability server sends its heartbeats to the **Heartbeat Transmit Port**.
- **Status Listen Port**—Double-Take Availability servers use the **Status Listen Port** to listen for requests from the Replication Console and other clients.
- **Time Between Heartbeats**—All Double-Take Availability servers transmit a heartbeat. This heartbeat allows other Double-Take Availability servers and Double-Take Availability clients to locate and identify the Double-Take Availability servers. The heartbeat is a broadcast UDP transmission. This heartbeat can be disabled, but if it is, Double-Take Availability will not auto-detect the Double-Take Availability servers to populate the Replication Console. By default, there are 3 seconds between heartbeats. If you set this option to 0, the heartbeats are disabled.
- **Missed Heartbeat Limit**—This is the number of heartbeats which can be missed before transmission is stopped and data is queued on the source.
- **Show Heartbeat Messages in Logger**—This checkbox enables the heartbeat messages in the Double-Take Availability log. Enabling this option will cause your logs to fill up faster.
- **TCP Window Size**—This option is the size, in bytes, of the buffer used for TCP transfers. This is an operating system buffer, not a Double-Take Availability buffer. If this option is set to zero (0), Linux kernel versions 2.6.7 or later can automatically tune this buffer setting for best server performance. Therefore, the recommended setting is 0 for automatic tuning, if you are using a version 2.6.7 or later Linux kernel. If you want to reduce or control network traffic, you can configure this option to a static size. The default is 375000 for a 1 GB network. Modifications should be relative to that speed using the calculation 37500 * network_speed_in_bits_per_second / 100 Mbit.



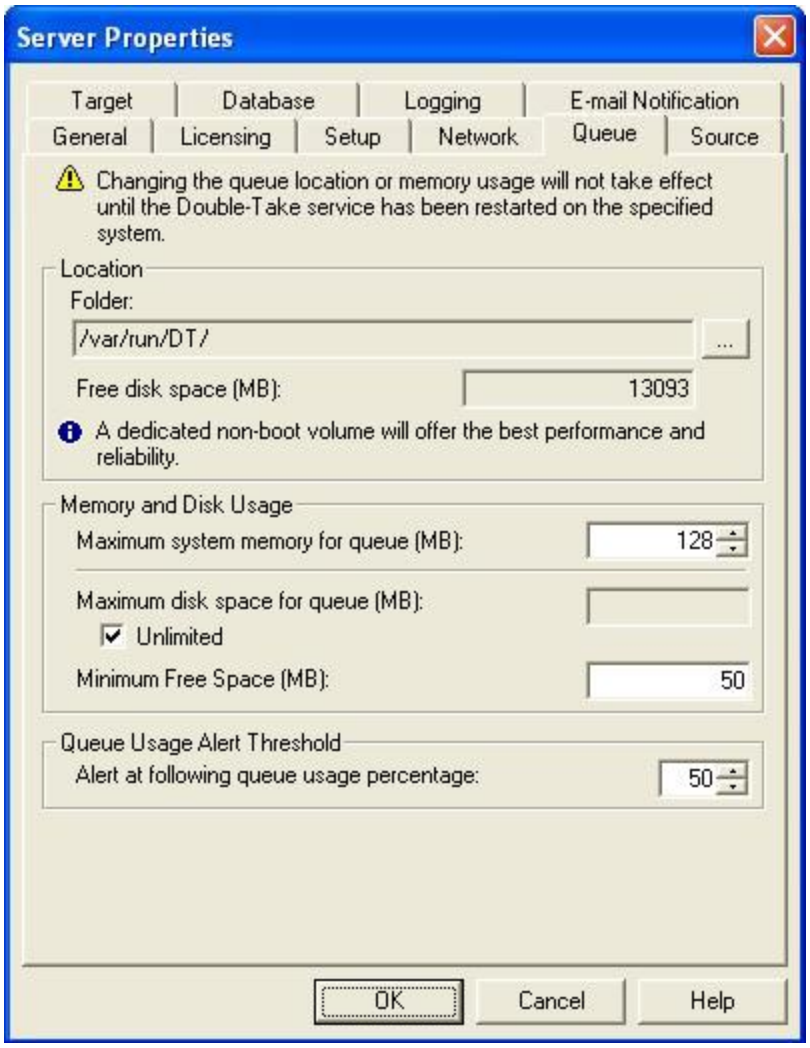
If you want to control network traffic, you may find the Double-Take Availability bandwidth limiting features to be a better method.

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

Queuing data

You should configure queuing on both the source and target.

1. 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 daemon has been restarted on the server.

  Select a location on a volume that will have minimal impact on the operating system and applications being protected. For best results and reliability, this

should be a dedicated, non-boot volume. The disk queue should not be on the same physical or logical volume as the data being replicated.



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 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 of memory is used. 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 daemon 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**.



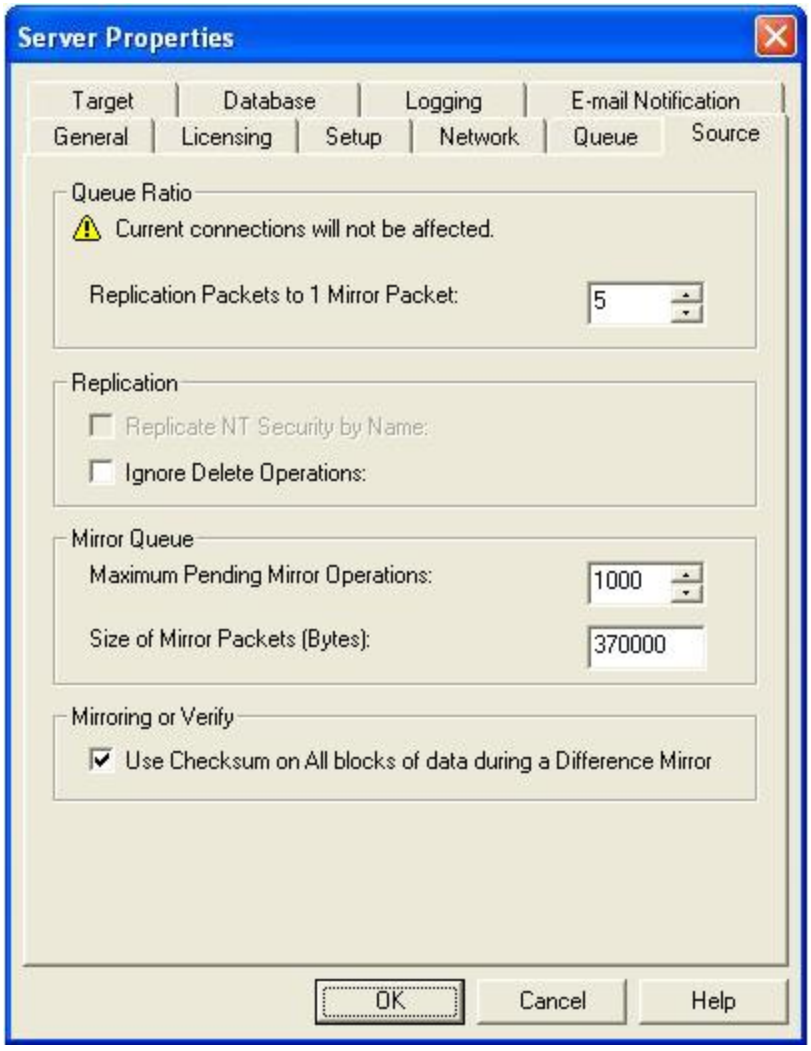
The **Maximum disk space for queue** and **Minimum Free Space** settings work in conjunction with each other. For example, assume your queues are stored on a 10 GB disk with the **Maximum disk space** for queue set to 10 GB and the **Minimum Free Space** set to 500 MB. If another program uses 5 GB, Double-Take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

- Specify the **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. Right-click a 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**—This is a Windows option only.
   - **Ignore Delete Operations**—This option allows you to keep files on the

target machine after they are deleted on the source. When a file is deleted on the source, that delete operation is not sent to the target. (All edits to files on the source are still replicated to the target; only deletions of whole files are ignored.) This option may be useful to give you an opportunity to make a backup of these files in the event they are needed in the future.



If delete operations are ignored long enough, the potential exists for the target to run out of space. In that case, you can manually delete files from the target to free space.

- **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 32768 bytes.
- **Use Checksum on All blocks of data during 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.

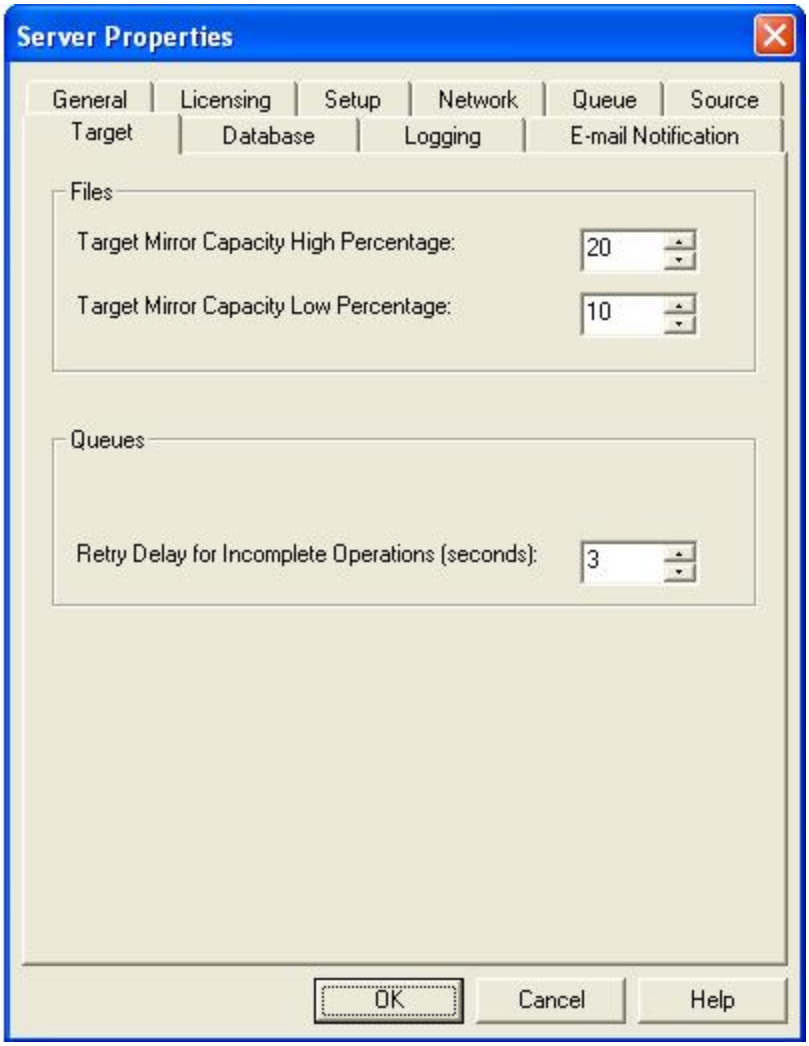


Database applications may update files without changing the date, time, or file size. Therefore, if you are using database applications, you should use the Block Checksum All option to ensure proper file comparisons.

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

Configuring target data processing options

1. Right-click a server on the left pane of the Replication Console.
2. Select **Properties**
3. Select the **Target** tab.



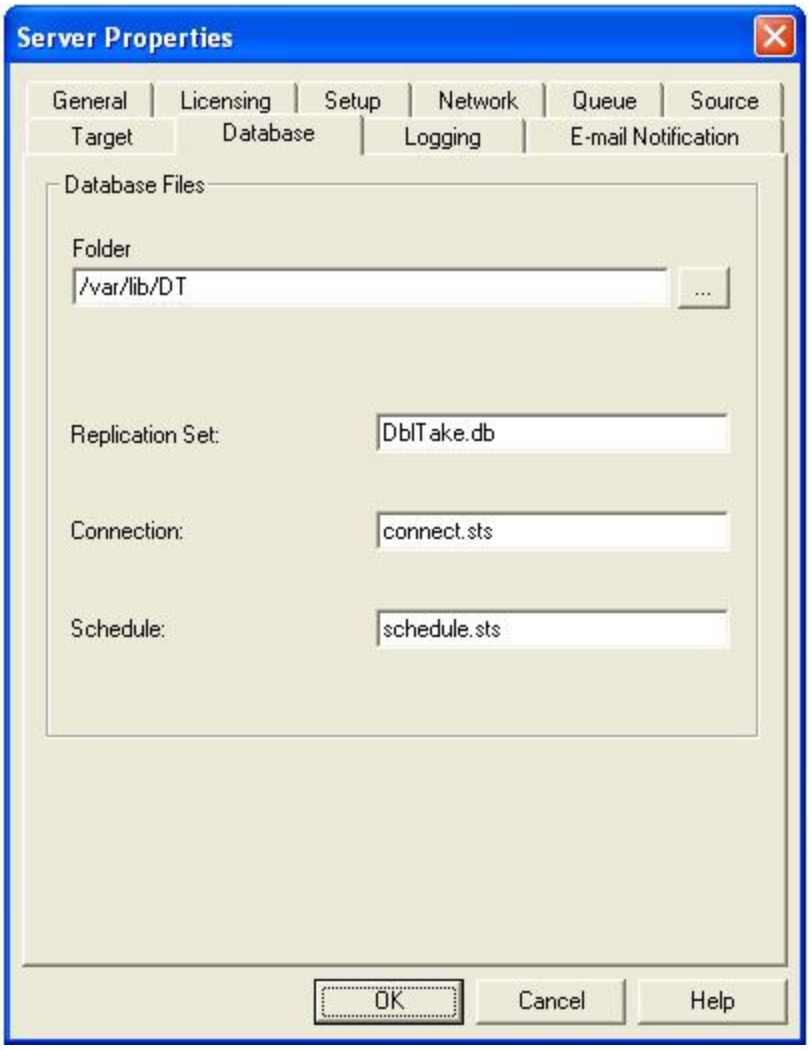
4. Specify how the target will process data.
   - **Target Mirror Capacity High Percentage**—You can specify the maximum percentage of system memory that can contain mirror data before the target signals the source to pause the sending of mirror operations. The default setting is 20.
   - **Target Mirror Capacity Low Percentage**—You can specify the minimum percentage of system memory that can contain mirror data before the target signals the source to resume the sending of mirror operations. The default setting is 10.

- **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.

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

Specifying the Double-Take Availability database storage files

1. Right-click a server on the left pane of the Replication Console.
2. Select **Properties**
3. Select the **Database** tab.



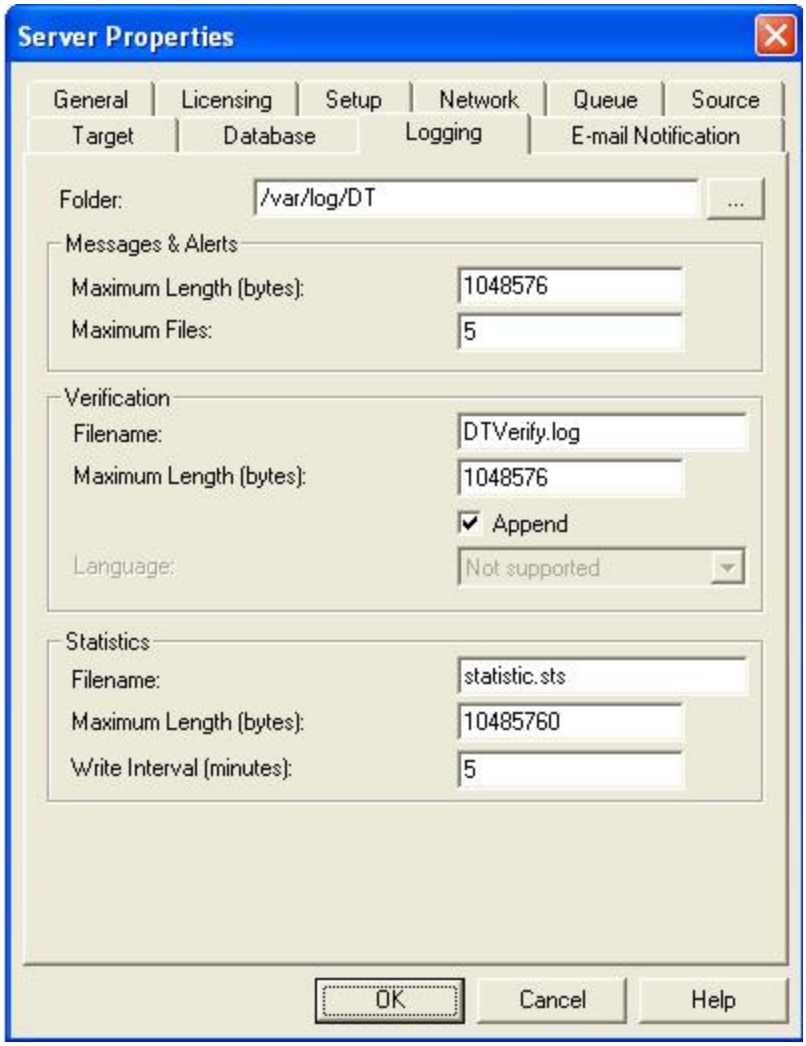
4. Specify the database files that store the Double-Take Availability replication set, connection, and scheduling information.
   - **Folder**—Specify the directory where each of the database files on this tab are stored. The default location is the directory where the Double-Take Availability program 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. Right-click a 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 daemon log files. The default size is 1048576 bytes and is limited by the available hard drive space.

- **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 prepended 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.
  - **Maximum Length**—Specify the maximum length of the verification log file. The default maximum length is 1048576 bytes (1 MB).
  - **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.
  - **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.
  - **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.
  - **Write Interval**—Specify how often Double-Take Availability writes to the statistics log file. The default is every 5 minutes.

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

E-mailing system messages

You can e-mail system messages to specified 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. To enable e-mail notification for a server, right-click the server in the left pane of the Replication Console and select **Properties**.
2. Select the **E-mail Notification** tab.
3. Select **Enable notification**.

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

4. Specify your e-mail settings.
   - **Mail Server (SMTP)**—Specify the name of your SMTP mail server.

     Specifying an SMTP server is the preferred method because it provides a direct connection between the mail server and Double-Take Availability, which decreases message latency and allows for better logging when the mail server cannot be reached.

     If you do not specify an SMTP server, Double-Take Availability will attempt to use the Linux mail command. The success will depend on how the local mail system is configured. Double-Take Availability will be able to reach any address that the mail command can reach.

   - **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.



You can test e-mail notification by specifying the options on the E-mail Notification tab and clicking **Test**. 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.

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 daemon, 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.

Chapter 18 Security

To ensure protection of your data, Double-Take Availability offers multi-level security using native operating system security features. Privileges are granted through membership in user groups. The groups can be local or LDAP (Lightweight Directory Access Protocol). To gain access to a particular Double-Take Availability source or target, the user must provide a valid local user account that is a member of one of the Double-Take Availability security groups. Once a valid user name and password have 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. This security group name is dtamin, and the default group ID is 501.
- **Monitor Access**—Servers and statistics can be viewed, but functionality is not available. This security group name is dtmon, and the default group ID is 502.
- **No Access**—Servers appear in the clients, but no access to view the server details is available.

Although Double-Take Availability passwords are encrypted when they are stored, Double-Take Availabilitysecurity design does assume that any machine running the Double-Take Availability client application is protected from unauthorized access. If you are running the Double-Take Availability client and step away from your machine, you must protect your machine from unauthorized access.

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 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.

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.

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**.

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**.

Chapter 19 Evaluating Double-Take Availability

The following evaluation procedure has eleven tasks containing step-by-step instructions for evaluating the core functionality of Double-Take Availability, specifically mirroring, replication, failover, and restoration. This is a good process for users who want to see, first-hand, the benefits that Double-Take Availability has to offer.

Before starting this evaluation procedure, make sure you have Double-Take Availability installed on the source and target. You should have at least 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 even of a source failure, the application can be started on the target and the files can then be accessed.

This evaluation consists of the following tasks.

- Establishing a connection
- Monitoring the activity and completion of the initial mirror
- Changing data to cause replication
- Verifying the data changes on the target
- Testing your target data
- Configuring failover monitoring
- Monitoring failover
- Simulating a failure
- Simulating data changes after failover
- Initiating failback
- Restoring your data

Establishing a connection

1. Start the Double-Take Availability Replication Console by selecting **Start**, **Programs**, **Double-Take for Linux**, **Double-Take 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**.

   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.

   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.

   Double-Take Availability will automatically attempt to log on to the selected source using previously cached credentials. If the logon is not successful, the Logon dialog box will appear prompting for your security identification.

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.

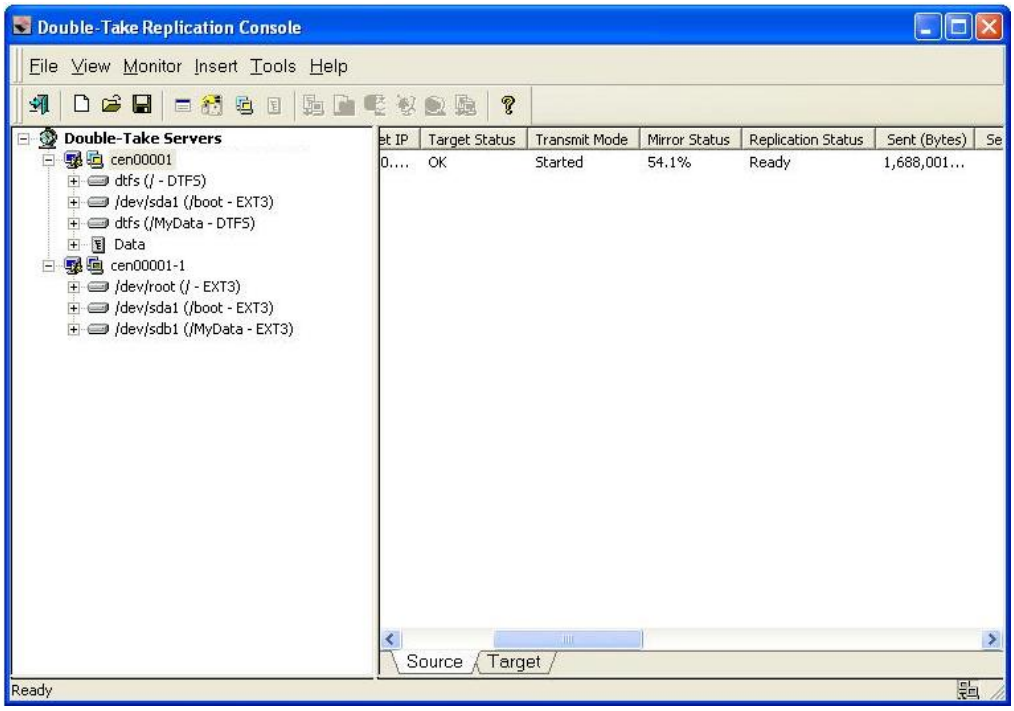
   Double-Take Availability will automatically attempt to log on to the selected targetusing previously cached credentials. If the logon is not successful, the Logon dialog box will appear prompting for your security identification.

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. For example, /var/data is transmitted to /var/data. Click **Next** to continue.
10. 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.
11. 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 various 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 of the Replication Console window 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 connection through the Replication Console 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 connection through the Replication Console contains complete details on all of the Replication Console statistics.



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.

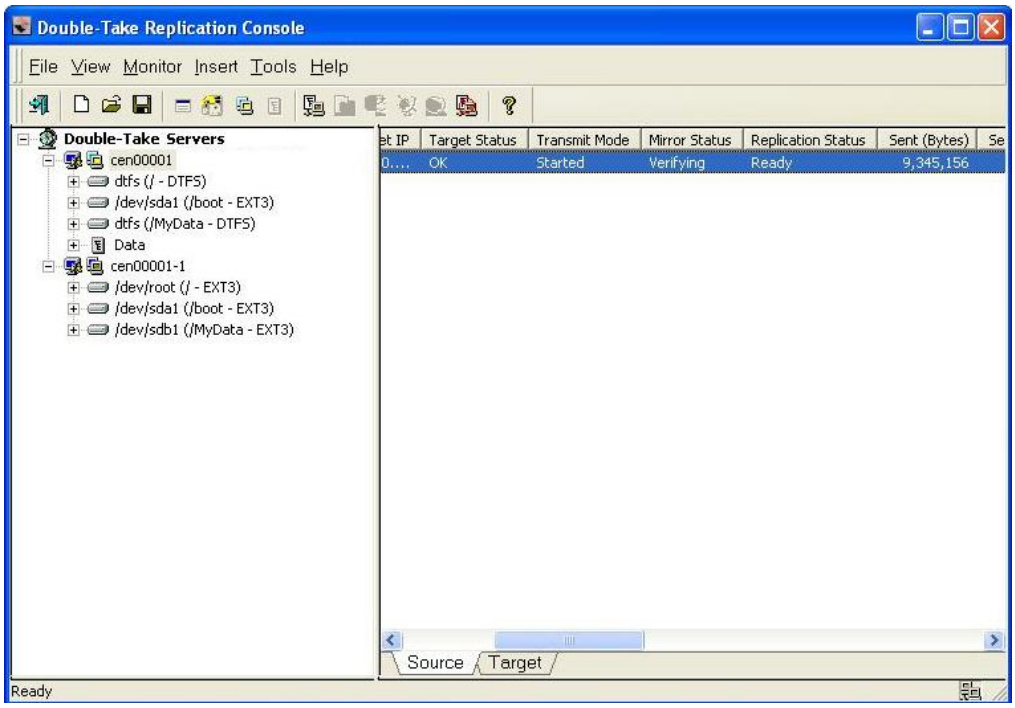
Verifying the data changes on the target

Now that you have modified some of the files, you want to be sure that the file modifications were applied correctly.



Machines that are doing minimal or light processing, as you are in this evaluation, may have file operations that remain in memory until additional operations flush them out and they are written to disk. This may make Double-Take Availability files on the target appear as if they are not synchronized. When the operations in memory are released, the files will be updated on the target. To make sure this does not impact your testing, flush memory by copying a couple of files from one directory to another and then deleting them. You can also use the sync program (which exercises the sync(2) system call) to flush memory.

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 options 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 data to the target automatically**.

Since your target file is newer, make sure that **Only if the source's date is newer than the target's date** 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.

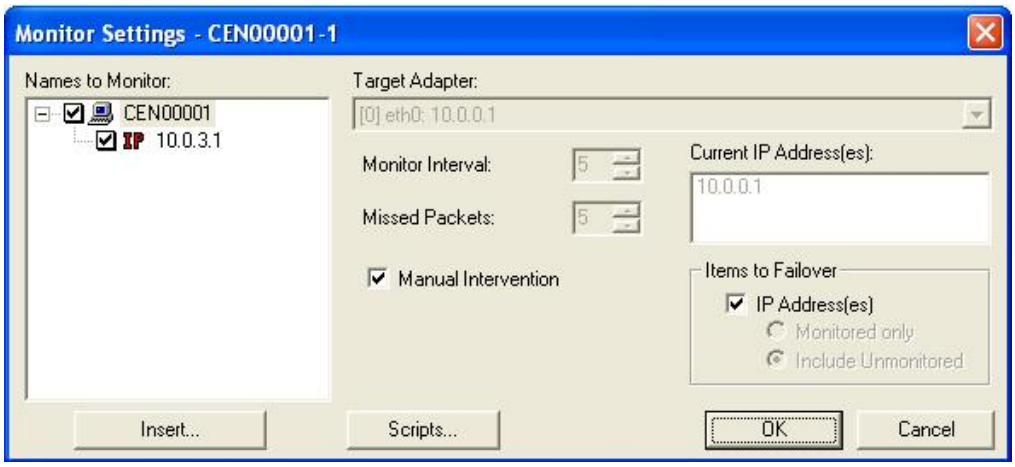
- **User data**—If you are protecting user files, you can use the associated application to open the files on the target. Open one or more of the files to test the integrity of the data. Do not save the file after you have opened it, because that will update the copy of the data on the target, which 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. In order 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. This will give 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 **Difference Mirror**. Make sure that **Only send if source's date is newer than the target's date** 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. The Failover Control Center can be started from within the Replication Console or from the Windows desktop.
   - From the Replication Console, select **Tools**, **Failover Control Center**.
   - From the Windows desktop, select **Start**, **Programs**, **Double-Take for Linux**, **Availability**, **Double-Take 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.

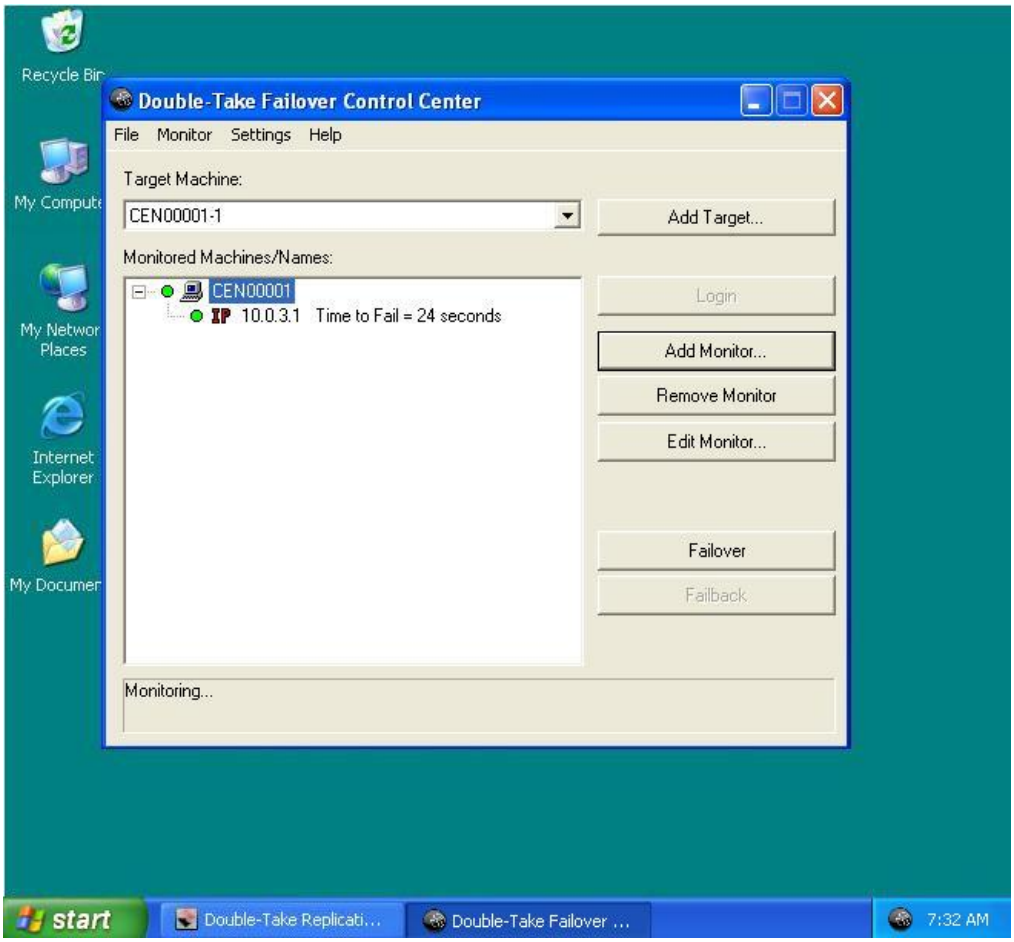


At this point, in terms of your evaluation, your failover configuration is complete because you will be using the default settings for the remaining options. But while you are viewing the Monitor Settings dialog box, notice the configuration options available to you.

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



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 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.



The Linux system log on the target provides details on the actual steps that have occurred during failover.

4. Ping the source's IP address from a client machine.
5. 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.

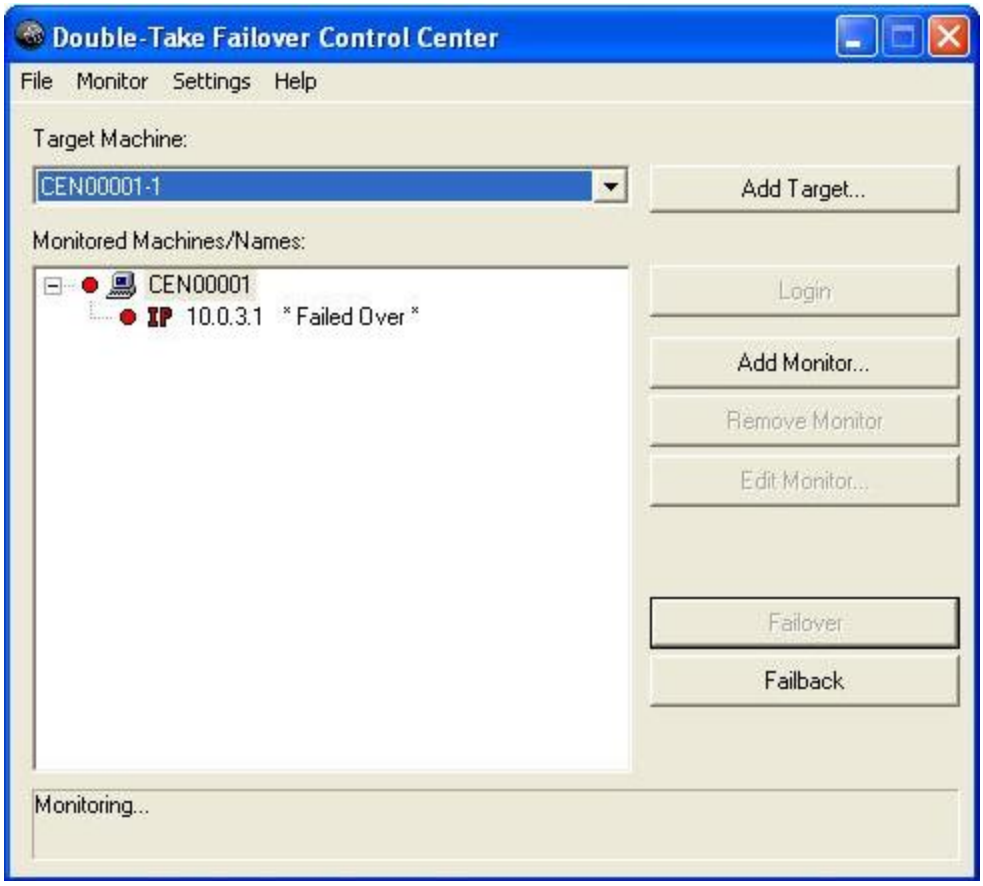
Initiating failback

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.

   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.

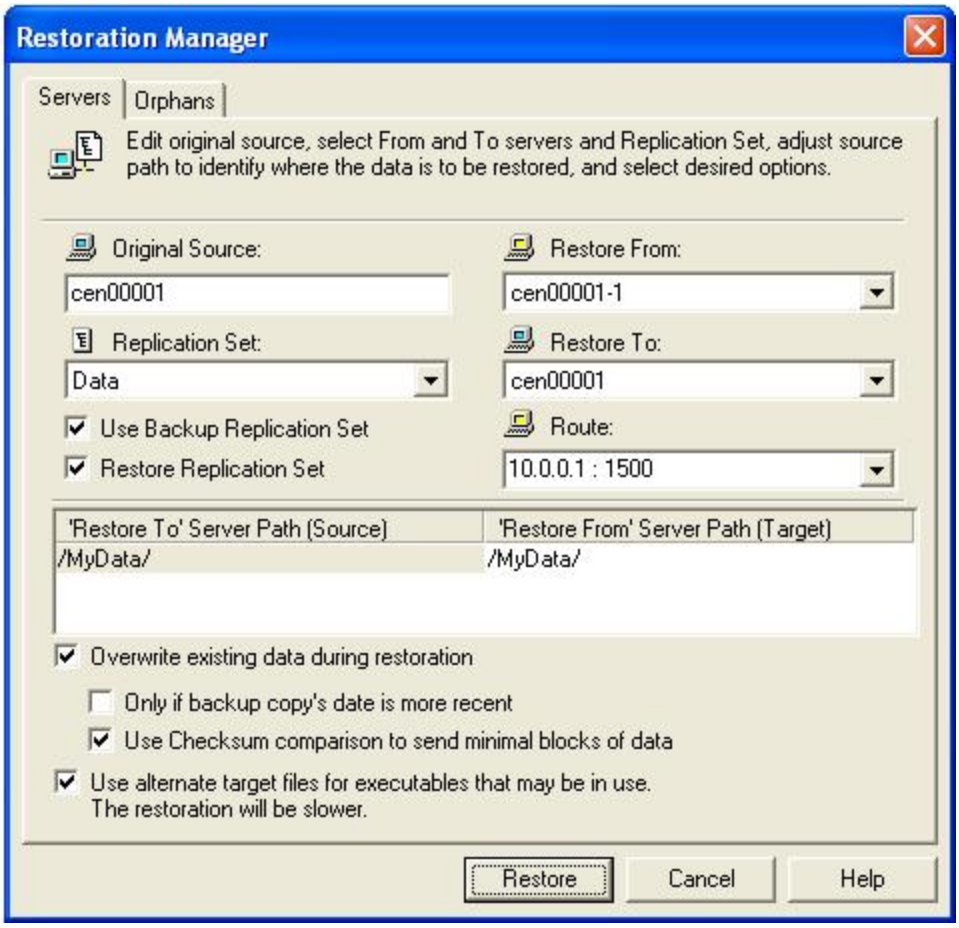
6. At this time, you would connect the source to the network. For this evaluation, reconnect the network cable(s) on the source that you disconnected to simulate the failure.
7. After the source is online, select **Stop** in the Failover Control Center to indicate that you do not want to continue monitoring 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.

Once the restoration is complete, your evaluation is complete. Congratulations!

Notices

© 1996-2011 Double-Take Software, Inc. All rights reserved.

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. Linux is a registered trademark of Linus Torvalds. Red Hat is a registered trademark of Red Hat, Inc. All other trademarks are the property of their respective companies.



Corporate Headquarters
Vision Solutions, Inc.
Irvine, California USA
Tel: +1 (949) 253-6500

Index

.

.rpm 24

A

activation code 41, 211
auto-disconnect 125
auto-reconnect 125, 127
auto-remirror 53, 133

B

bandwidth limiting 167
block device
  adding 50, 144, 148
  configuration 41

C

chained configuration 12
compression 169
configurations
  chained 12
  many-to-one 12
  one-to-many 12
  one-to-one 12
connection 47, 53, 62
  database storage file 225
  disconnect 129
  firewall 58
  ID 119
  NAT 58
  overview 119
  reconnect 127
  simulation 61
  target processing 128
  types 46
Connection Manager 53
Connection Wizard 47, 62

D

daemon 41
data
  failover monitoring 180
  monitoring 66
data loss 6
data protection 46
database files 225

disconnect connection 129
documentation 41
dtfs_mounts 41
dtloop 41
DTSetup 41

E

e-mail notification 98, 229
erasing 27

F

failback
  overview 171
failover
  configuring failover monitoring 172
  editing failover monitoring configuration 179
  failover 183
  monitoring 180
  overview 7, 171
  removing failover monitoring configuration 183
Failover Control Center 42, 45
  refresh rate 45
file differences mirror 53, 133
  options compared 131
file system 17
  configuration 41
firewall 58
full-server protection 11
full-server restoration 198, 202

H

hard links 138

I

ICMP 58
identity 209
installation
  Linux servers 24
  overview 22
  Windows client 26

K

kernel type 17
kernel version 17

L

licensing 211
logging
  e-mailing system messages 98, 229
  log file 73, 77, 227
  messages 77
  verification 159
logging on and off
  Replication Console 233

M

manual intervention 172, 183
many-to-one configuration 12
mirroring 53
  automatically 133
  controls 131
  overview 7, 130
monitoring
  data workload 66
  failover data 180
monitoring tools
  overview 65
  SNMP 111

N

named pipes 233
NAT 58
network communications 216

O

one-to-many configuration 12
one-to-one configuration 12
orphan files 135
overview 6, 12

P

ports 41, 58, 216
protection 46

Q

queues 125
  overview 120
  queuing data 122, 218

R

reconnecting automatically 127
recurse subdirectories 142, 146
remirror 53
removing 27
replication
  capabilities 138
  configuration 41
  inserting tasks 153
  overview 7, 137
  starting 152
Replication Console 42-43
  credentials 44
  logging on and off 233
  workspaces 43
replication set
  block device 148
  calculating size 150
  copying 149
  creating 50, 144
  creating manually 146
  database storage file 225
  deleting 151
  limitiations 142
  modifying 149
  overview 142
  renaming 149
  rules 146
requirements 17
restoration
  overview 7

S

security
  credentials 44
  groups 41
  overview 232
server identity 209
server settings 208
SNMP
  configuration 111
  overview 111
  statistics 115
  traps 112
soft links 138

source 6
  data processing options 221
  server startup options 214
statistics
  file 101-102, 227
  output 103
  overview 101
  SNMP 115
synchronization 7
system messages
  e-mail 98, 229

T

target 6
  data procesing options 223
  pause 128
  resume 128
task command processing 153
TDU 61
transmission
  bandwidth limiting 167
  controls 161
  network communications 216
  overview 160
  schedule database storage file 225
  scheduling 161

U

UDP 58
uninstallation 27
upgrade
  Linux servers 24
  overview 22

V

verification
  log file 159, 227
  manually verifying 155
  overview 154
  scheduling 157
virtual systems 17, 202

W

workspaces 43