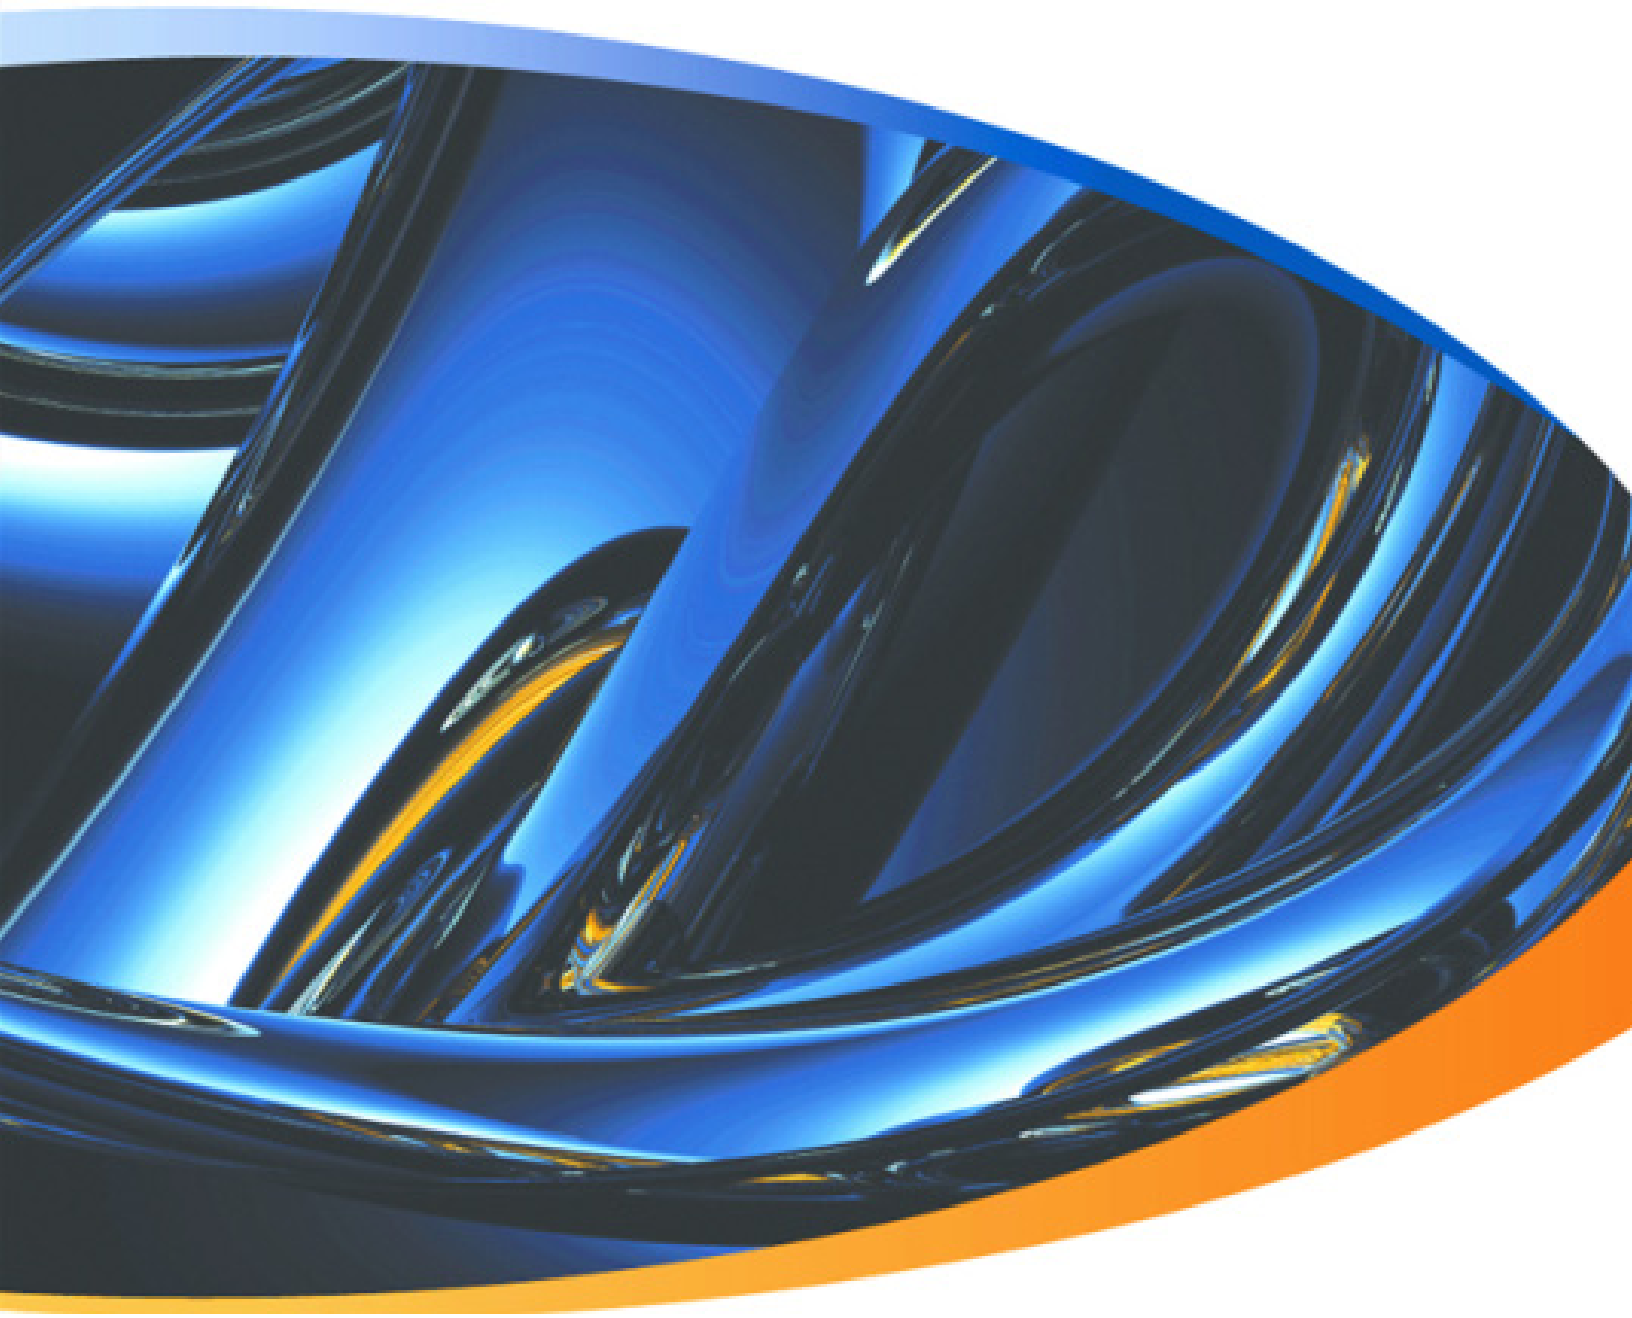
Double-Take for Linux

Double-Take for Linux User's Guide



Double-Take, Balance, 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. in the United States and other countries. Novell, the Novell logo, the N logo, SUSE are registered trademarks of Novell, Inc. in the United States and other countries. All other trademarks are the property of their respective companies.

© 1996-2009 Double-Take Software, Inc. All rights reserved.

Version 4.5 DTL-45-UG-040809

Attention Notice

IMPORTANT: READ CAREFULLY BEFORE USING THIS PRODUCT AND ACCOMPANYING DOCUMENTATION ("Software"). USE OF THIS SOFTWARE IS SUBJECT TO THE DOUBLE-TAKE SOFTWARE, INC. ("Double-Take") SOFTWARE LICENSE TERMS AND CONDITIONS SET FORTH BELOW ("License Agreement"). OPENING, INSTALLING, AND/OR USING THE SOFTWARE INDICATES ACCEPTANCE OF THIS LICENSE AGREEMENT. IF YOU ("Licensee") DO NOT ACCEPT THE TERMS AND CONDITIONS IN THIS LICENSE AGREEMENT, LICENSEE MUST NOT INSTALL AND/OR USE THE SOFTWARE.

A. Controlling Language

All documents contained herein are in English and may be translated to different languages. In the event of any inconsistencies between the English and translated versions of any document, the English Version shall control. Copies of the English and translated documents can be obtained as follows: (1) from the physical media; (2) downloaded from the Double-Take Online Support website at http://support.doubletake.com; or (3) by calling Double-Take Technical Support at the number listed on the Support website.

B. License Grant

Subject to Licensee's compliance with the terms and conditions set forth herein, in return for the fee paid by Licensee, Double-Take hereby grants Licensee a limited, non-transferable personal license to use the object code version of the Software included on the media on a single server authorized in writing by Double-Take in the relevant documentation pursuant to the terms and conditions specified by Double-Take in its License Agreement, quotation, invoice or other terms that may accompany the Software. If Your purchase includes multiple instances of Software, usage is restricted to the number of licenses specified in the relevant documentation accompanying Licensee's purchase. If Licensee has purchased a Gold Disk license, usage is restricted to the number of servers and licenses set forth in Licensee's purchase order or Gold Disk Agreement. If Licensee receives a license for evaluation purposes, usage is restricted to a controlled environment for evaluation and testing purposes only and solely for the period of time specified by Double-Take. The evaluation Software will be equipped with a time out mechanism set for the period of time specified by Double-Take. Evaluation licenses are not eligible for Indemnification.

Unless otherwise specified, all Software licenses will be perpetual unless terminated or transferred in accordance with the terms of this License Agreement.

In the event that the Software is Linux based, specific components of the Software are subject to additional licenses. Kernel dtfs.ko and dtrep.ko ("Open Source Software") are subject to the GNU General Public License version 2 ("GPL"), and LibESMTP is subject to the GNU Lesser General Public License Version 2.1 or later ("LGPL"), as published by the Free Software Foundation. The Open Source Software shall not be subject to the remaining terms herein. There are no warranties provided with respect to the Open Source Software and all implied warranties are disclaimed, in accordance with the terms of the GPL and/or LGPL. In the event of any conflict between the terms herein and the GPL and/or LGPL with respect to the Open Source Software, the terms of the GPL and/or LGPL shall control. A copy of the GPL and LGPL licenses are provided with the Software distribution and can also be obtained at http://www.gnu.org/licenses/. The source files for the Open Source Software are available via an FTP site from Double-Take, and can be obtained upon request submitted to LinuxPM@doubletake.com. The version of LibESMTP library file being utilized is available at http://www.stafford.uklinux.net/libesmtp/index.html. Notwithstanding the foregoing, the GPL and LGPL shall only apply to the Open Source Software and shall not apply, in whole or part, to the remaining Software, and in no event shall Licensee have any right to access or obtain the source code for the remaining Software.

C. Copyright

Software is owned and copyrighted by Double-Take. This License Agreement confers no title or ownership and is not a sale of any rights in the Software. Licensee agrees that this Software remains the property of Double-Take and that Licensee only has a license to use this Software. Double-Take reserves all other rights to the Software.

D. Restrictions

Licensee may not (1) rent, lease, sublicense, time-share, sell or otherwise transfer the Software or documentation except as expressly authorized in these terms; (2) make the Software available over the internet or any other publicly accessible network or technology; (3) remove any copyright, trademark, or other proprietary notices from the Software or the media; (4) make any copies of the documentation; (5) copy the Software except as expressly provided for herein; (6) assign any rights or obligations hereunder; or (7) reverse engineer, decompile, disassemble, or otherwise attempt directly or indirectly to discover, use, disclose or transfer any source code or other confidential information contained in the Software.

E. Copying

Licensee may make and maintain one backup copy of the Software provided it is used only for Licensee's own backup purposes and Licensee keep possession of all backup copies.

F. Pre-Printed Terms

Any pre-printed terms and conditions of any documents used by any third party in connection with this License Agreement shall not be binding on Double-Take and shall not be deemed to modify this License Agreement.

G. Termination

Double-Take may terminate Licensee's license upon notice for failure to comply with any of the license terms set forth herein. Immediately upon termination, the Software, and all copies of the Software, at Double-Take's option, will be destroyed or returned to Double-Take.

H. Limitation of Liability

EXCEPT TO THE EXTENT PROHIBITED BY LOCAL LAW, IN NO EVENT WILL DOUBLE-TAKE OR ITS SUBSIDIARIES, AFFILIATES, DIRECTORS, OFFICERS, EMPLOYEES, AGENTS OR SUPPLIERS BE LIABLE FOR DIRECT, SPECIAL, INCIDENTAL, CONSEQUENTIAL, PUNITIVE, OR OTHER DAMAGES (INCLUDING LOST PROFIT, LOST DATA, OR DOWNTIME COSTS), ARISING OUT OF THE USE, INABILITY TO USE, OR THE RESULTS OF USE OF THE SOFTWARE, WHETHER BASED IN WARRANTY, CONTRACT, TORT OR OTHER LEGAL THEORY, AND WHETHER OR NOT DOUBLE-TAKE WAS ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. In any case, Double-Take's entire liability under any provision of this License Agreement shall be limited to amount received by Double-Take during the twelve (12) month period prior to the date the cause of action arose. Licensee's use of the Software is entirely at Licensee's own risk.

I. Limited Warranty

Double-Take warrants only that the Software media will be free of physical defects for a period of ninety (90) days from delivery. The entire liability of Double-Take and Licensee's exclusive remedy for Software that does not conform to this limited warranty shall be the repair or replacement of the defective media. This warranty and remedy are subject to Licensee returning the defective media during the warranty period to Double-Take or to the distributor whom Licensee obtained the Software.

J. Indemnification

Infringement. Double-Take will, at Double-Take's expense, indemnify, defend and hold harmless Licensee (including officers, directors, employees and agents), against any actual third party claims (including but not limited to reasonable attorneys' fees and costs) that the Software infringes, violates or misappropriates a patent, copyright, trademark, or other property right of a third party in the United States by the Software in the form provided by Double-Take to Licensee provided that (1) Licensee notify Double-Take promptly upon learning that the claim might be asserted; (2) Double-Take has sole control over the defense of the claim and any negotiations for its settlement or compromise; and (3) Licensee take no action that, in Double-Take's judgment, impairs Double-Take's defense of the claim.

This indemnification obligation shall be effective only if (1) Licensee has made all payments required by the terms of this License Agreement; (2) Licensee has given prompt notice of the claim; and (3) the infringement does not result from Licensee's modification of the Software or the Software's incorporation with other software, hardware or apparatus not supplied by Double-Take.

Remedy. If such a claim is made, Double-Take, at its own expense, may exercise any of the following remedies that are reasonably and commercially practicable (1) obtain for Licensee the right to continue to use the Software consistent with this License Agreement; (2) modify the Software so it is non-infringing and in compliance with this License Agreement; or (3) replace the Software with non-infringing software that complies with this License Agreement. If none of these options are available, Double-Take shall refund to Licensee the remaining value of the Software as amortized over a thirty-six (36) month life.

K. Disclaimer

TO THE EXTENT ALLOWED BY LAW, THIS SOFTWARE IS PROVIDED TO LICENSEE "AS IS" WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, WHETHER ORAL OR WRITTEN, EXPRESS OR IMPLIED. DOUBLE-TAKE SPECIFICALLY DISCLAIMS ANY IMPLIED WARRANTIES OR CONDITIONS OF MERCHANTABILITY, SATISFACTORY QUALITY, NON-INFRINGEMENT, TITLE, ACCURACY OF INFORMATIONAL CONTENT, AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE RESULTS AND PERFORMANCE OF THE SOFTWARE IS ASSUMED BY LICENSEE. NO ORAL OR WRITTEN INFORMATION OR ADVICE GIVEN BY DOUBLE-TAKE SHALL CREATE A WARRANTY OR AMEND THIS "AS IS" WARRANTY. Some jurisdictions do not allow exclusions of implied warranties or conditions, so the above exclusion may not apply to Licensee to the extent prohibited by such local laws. Licensee may have other rights that vary from country to country, state to state, or province to province.

L. Government

If the Software is licensed for use in the performance of a U.S. government prime contract or subcontract, Licensee agrees that, consistent with FAR 12.211 and 12.212, commercial computer software, computer software documentation and technical data for commercial items are licensed under vendor's standard commercial license.

M. Export Restrictions

Licensee acknowledges that this Software is subject to the export laws of the United States and agrees to comply at all times with such laws. This Software or any components, data, code or technology thereof may not be exported except in full compliance with all United States and other applicable laws and regulations. Licensee hereby represents and warrants that Licensee: (1) is not a citizen or resident of Cuba, Iraq, Libya, Sudan, North Korea, Iran, or Syria; (2) and if a legal entity, (a) is not an entity formed under the laws of Cuba, Iraq, Libya, Sudan, North Korea, Iran, or Syria; or (b) is not included on the U.S. Treasury Department's list of Specially Designated nationals or the U.S. Commerce Department's Table of Deny Orders.

N. Severability

If any term or provision herein is determined to be illegal or unenforceable, the validity or enforceability of the remainder of the terms or provisions herein will remain in full force and effect.

O. Gold Disk Maintenance Renewals

If Licensee chooses to renew maintenance, Licensee shall purchase maintenance for all licenses purchased hereunder and for any subsequent licenses added to the Gold Disk.

P. Notice

Any notices or other communications under this License Agreement shall be in writing and delivered in person, or a national overnight express delivery service or United States mail (registered or certified mail, postage pre-paid, return receipt requested) and shall be deemed to have been given upon confirmation of Double-Take's signature. Notices should be sent to the attention of Chief Financial Officer, Double-Take Software, Inc., 8470 Allison Pointe Blvd., Suite 300, Indianapolis, IN 46250 with a copy to Associate Corporate Counsel.

Rev 02/2009

© 1996-2009 Double-Take Software, Inc. All rights reserved.

Table of Contents

Chapter 1 Introduction 1-1

Double-Take operations 1-3

Mirroring 1-3
Replication 1-4
Failure monitoring and failover 1-5
Restoration 1-6

Sample configurations 1-7

One-to-one, active/standby 1-7
One-to-one, active/active 1-8
Many-to-one 1-9
One-to-many 1-10
Chained 1-11
Single server 1-12
Virtual systems 1-13

Resources 1-14

Chapter 2 Installation 2-1

System requirements 2-1

Installing or upgrading Double-Take for Linux 2-3

Configuring Your Double-Take Installation 2-4

Installing or Upgrading the Double-Take Client for Windows 2-9

Removing (Erasing) Double-Take 2-10

Removing the Double-Take Client for Windows 2-10

Chapter 3 Clients 3-1

Management Console 3-1

Logging on and off of Double-Take 3-2
Double-Take workspaces 3-3
Saving a workspace 3-3
Opening a workspace 3-3

Failover Control Center 3-4

Configuring communication ports 3-4
Configuring refresh rate 3-4
Clearing maintained security credentials 3-4

Chapter 4 Replication Sets 4-1

Replication capabilities 4-3

Creating a replication set 4-5

Creating or modifying replication rules manually 4-6

Selecting a block device for replication 4-7

Modifying a replication set 4-7

Renaming and copying a replication set 4-7

Deleting a replication set 4-7

Calculating replication set size 4-8

Chapter 5 Establishing a Connection Using the Connection Wizard 5-1
Chapter 6 Establishing a Connection Using the Connection Manager 6-1
Chapter 7 Establishing a Connection Across a NAT or Firewall 7-1
Verifying the Double-Take port settings 7-1
Configuring the hardware 7-2
Inserting the servers in the Management Console 7-2
Establishing the connection 7-2
Chapter 8 Simulating a Connection 8-1
Gathering system information 8-1
Establishing the simulated connection 8-2
Chapter 9 Managing Connections 9-1
Monitoring connections 9-2
Monitoring the message window 9-6
Queuing Double-Take data 9-8
Understanding queuing 9-8
Configuring queuing 9-10
Configuring automatic reconnections 9-12
Configuring compression 9-13
Disconnecting a connection 9-13
Chapter 10 Mirroring and Replication 10-1
Controlling mirroring 10-1
Controlling an automatic mirror 10-3
Configuring and removing orphan files 10-5
Starting replication 10-6
Chapter 11 Verification 11-1
Manual verification 11-1
Scheduled verification 11-2
Verification log 11-3
Chapter 12 Data Transmission 12-1
Controlling transmission manually 12-1
Controlling a target manually 12-1
Scheduling transmission criteria and bandwidth limiting 12-2

Chapter 13 Failover and Failback 13-1
Configuring failover 13-1
Editing failover configuration 13-3
Removing failover configuration 13-4
Monitoring failover 13-4
Testing failover 13-5
Failover manual intervention 13-5
Initiating failback 13-6
Chapter 14 Restoration 14-1
Chapter 15 Monitoring Tools 15-1
Log files 15-2
Viewing the log files 15-3
Log messages 15-4
Linux System Log 15-9
DTStat 15-19
Configuring the statistics file 15-20
Running DTStat from the command prompt 15-21
DTStat statistics 15-22
SNMP 15-26
Configuring SNMP 15-26
SNMP statistics 15-27
SNMP traps 15-29
Chapter 16 Server Settings 16-1
General 16-2
Setup 16-4
Network 16-6
Queue 16-7
Source 16-9
Target 16-11
Database 16-12
Logging 16-13
E-mail notification 16-14
Chapter 17 Security 17-1
Security access levels 17-1
Security advantages and considerations 17-1
Clearing maintained security credentials 17-2

1 Introduction

Welcome to Double-Take® for Linux®! Double-Take is a real-time data replication and failover software product. Double-Take 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 allows you to specify mission-critical data that must be protected and replicates, in real-time, that 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 critical data from the source. Double-Take monitors any changes to the critical data and sends the changes to the target server. By replicating only the file changes rather than copying an entire file, Double-Take allows you to more efficiently use resources.

- **Offsite disaster recovery services**—No business is immune from the many disasters - disk crashes, power failures, human error, natural disasters - that will inevitably stop the flow of data at one or more of your facilities. Tape-based disaster recovery can only restore data to the point of the last backup, which was most likely the prior night. Any data created since the last backup will be lost. An effective disaster recovery plan requires a comprehensive data protection plan, including Double-Take continuous data replication.

  Double-Take can be combined with your existing tape backup solution for a more comprehensive disaster recovery plan. Double-Take is a disaster recovery software based on asynchronous real-time replication and automatic failover to provide cost-effective business continuity for Microsoft® Exchange, Microsoft SQL Server®, Oracle®, virtual systems, file servers, and many other applications. Double-Take provides continuous data protection by sending an up-to-the-minute copy of the data as it is being changed to the target server. Features such as built-in bandwidth control allows data to be replicated to a remote source, far from harms way of a disaster such as hurricanes, tornadoes, and brown-outs. That same bandwidth control allows data to be restored within minutes.

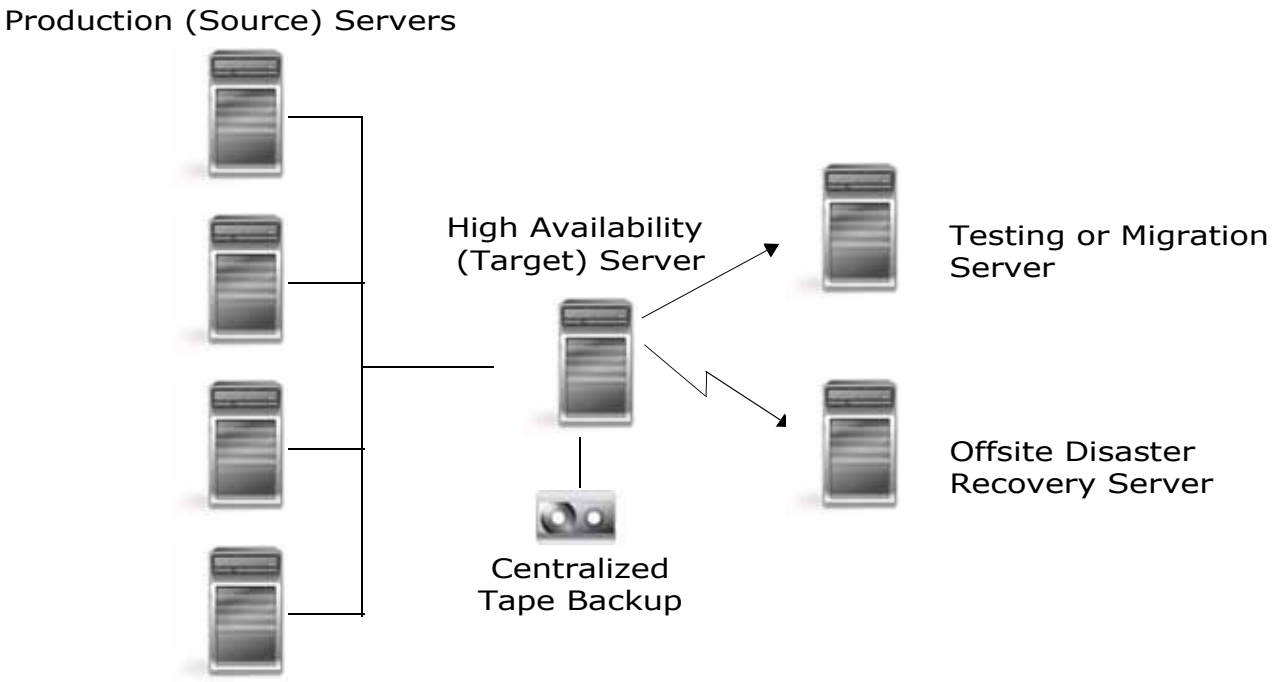
- **Local high availability services**—As businesses come to increasingly depend on continuous access to their data, ensuring that their data is available on-demand is of paramount importance. Traditional solutions, such as tape backup and hardware mirroring, are not without flaws. Tape backup poses a potential risk in that data backups must be performed when the system is idle, meaning that as much as a day's worth of data could be lost. Tape recovery time can also be below many business's recovery time objectives (and does not provide high availability).

  A more cost-effective answer to a high availability solution is asynchronous host-based replication for high availability. Support for non-proprietary hardware and storage systems means you can leverage your existing resources. You will also get real-time data protection without distance limitations, ensuring high availability for all your business-critical applications, including e-mail. Real-time data replication at the byte level, regardless of application, is a much more efficient use of computing and bandwidth resources for high availability. With Double-Take all of your applications can have cost-effective, real-time data protection and high availability.

- **Enhanced centralized backup**—The rapid growth in storage brought on by the Internet and distributed computing has placed nearly impossible demands on administrators responsible for protecting corporate data assets. The backup window has shrunk to nearly zero and tape backup systems can introduce significant overhead to a production server, seriously impacting its performance. While the importance of backups increases, the impact of periodic full system backups is obvious. Even nightly incremental backups dominate processing while they examine every file system object and then read all files that have changed in their entirety for backup. Performing this process across a network adds additional overhead as the entire process happens across the wire.

  These days permanent point in time storage and recovery, like that provided by periodic tape backup, is required. And despite the fact that Double-Take cannot provide a way to retrieve historical file versions or files that may have been previously deleted by users, Double-Take can enhance the backup process by continuously replicating critical data to centralized servers and using tape backup systems to backup the replica rather than the production servers. Using Double-Take offloads the burden of periodic tape backups from multiple production servers to a dedicated backup server and makes centralized tape backup a reality, significantly reducing management cost and improving reliability. Regardless of a file’s state on the source, on the target every file is closed and available for consistent backup at any point in time.

The following diagram is one common Double-Take configuration.



Double-Take operations

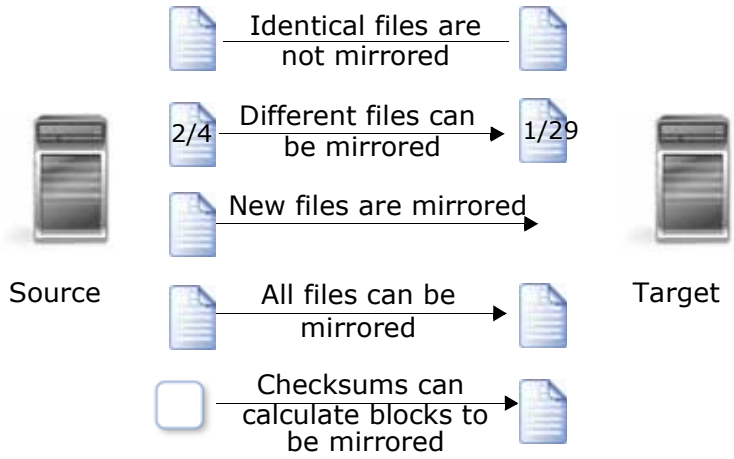
Double-Take performs four basic types of operations.

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

Mirroring

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

If subsequent mirroring operations are necessary, Double-Take 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.



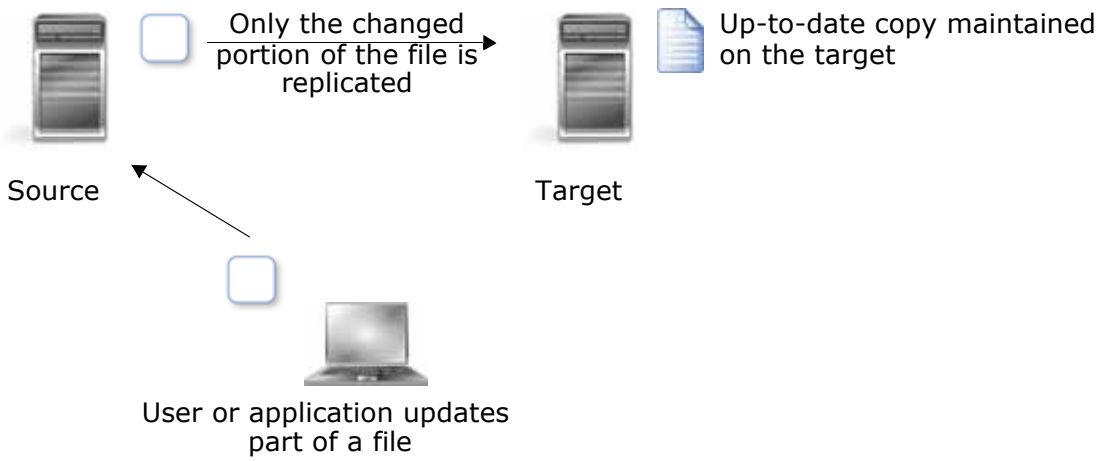
Flexible mirroring options allow you to choose which files are mirrored from the source to the target.

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

Replication

Replication is the real-time transmission of file changes. Unlike other related technologies, which are based on a disk driver or a specific application, the Double-Take 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, file system changes from the user-defined data set are tracked. Double-Take immediately transmits these file changes to the target server. This real-time replication keeps the data on the target up-to-date with the source and provides high availability and disaster recovery with minimal data loss.



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

Failure monitoring and failover

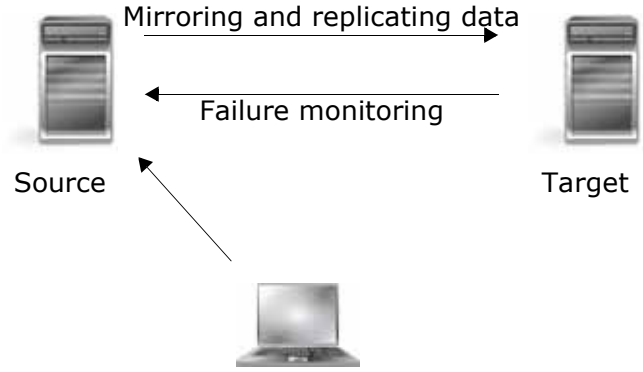
Failover is the process in which a target stands in for a failed source. As a result, user and application requests that are directed to the failed source are routed to the target.

Double-Take 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 assumes that the server has failed. Double-Take then prompts the network administrator to initiate failover, or, if configured, it occurs automatically.

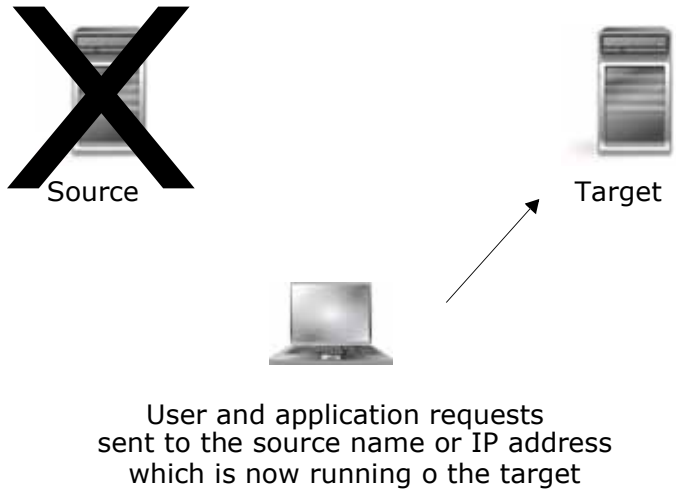
The failover target assumes the network identity of the failed source. When the target assumes the identity of the source, user and application requests destined for the source server or its IP address(es) are routed to the target.

When partnered with the Double-Take 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.

Before failover occurs



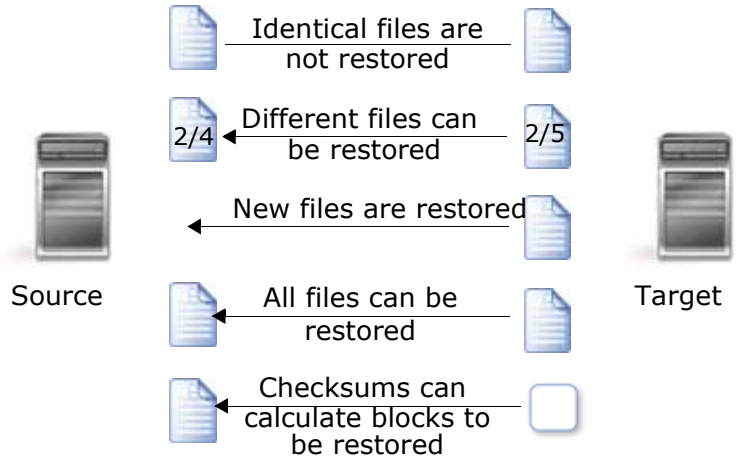
After failover occurs



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.

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 target will contain the same data as your Double-Take source. If you are using the Double-Take 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.



Flexible restoration options allow you to choose which files are restored from the target to the source.

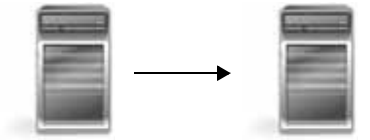
When a restoration is complete, the source and target are again synchronized.

Sample configurations

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

NOTE: The Double-Take replication and failover features have different server requirements. If you use Double-Take to monitor and stand in for failed servers, the source and target must have identical versions of operating system and applications. By contrast, Double-Take replication operations do not require that the source and target servers have identical versions of operating system and applications.

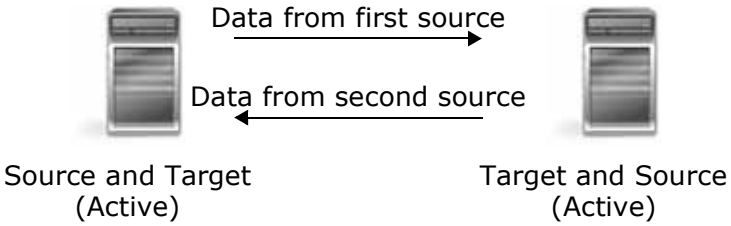
One-to-one, active/standby



Source (Active) Target (Standby)

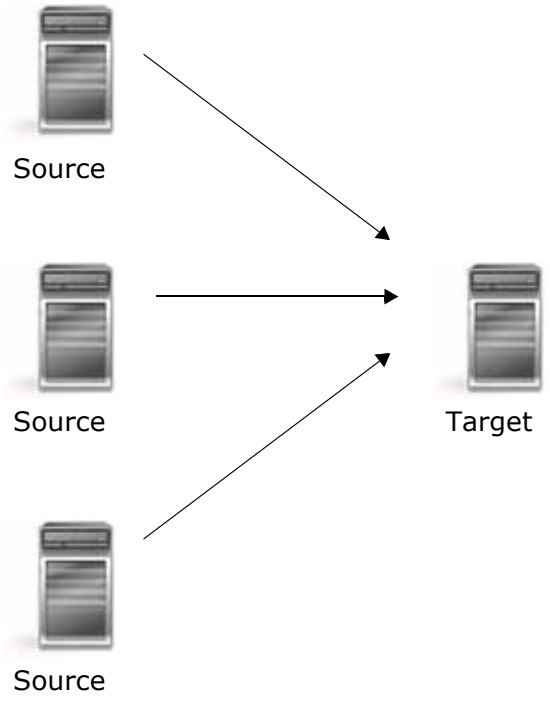
Description	One target server, having no production activity, is dedicated to support one source server. The source is the only server actively replicating data.
Applications	This configuration is appropriate for offsite disaster recovery, failover, and critical data backup. This is especially appropriate 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.

One-to-one, active/active



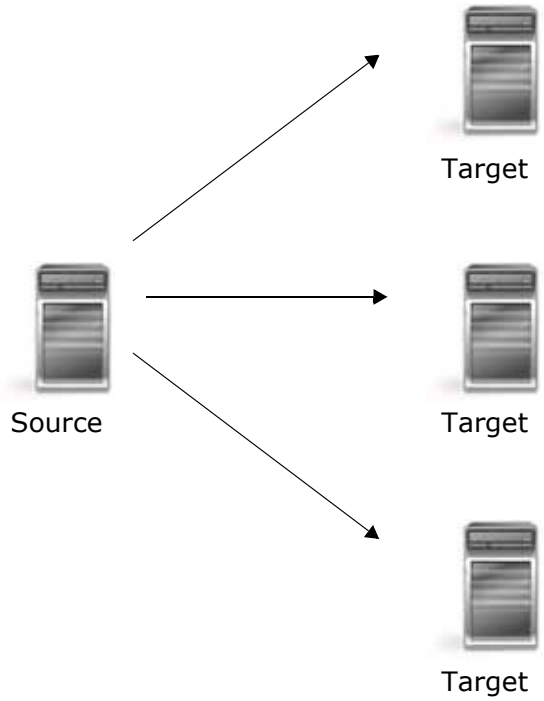
Description	Each server acts as both a source and target actively replicating data to each other.
Applications	This configuration is appropriate for failover and critical data backup. This configuration is more cost-effective than the Active/Standby configuration because there is no need to buy a dedicated target server for each source. In this case, both servers can do full-time production work.
Considerations	Coordination of the configuration of Double-Take 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.

Many-to-one



Description	Many source servers are protected by one target server.
Application	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.

One-to-many



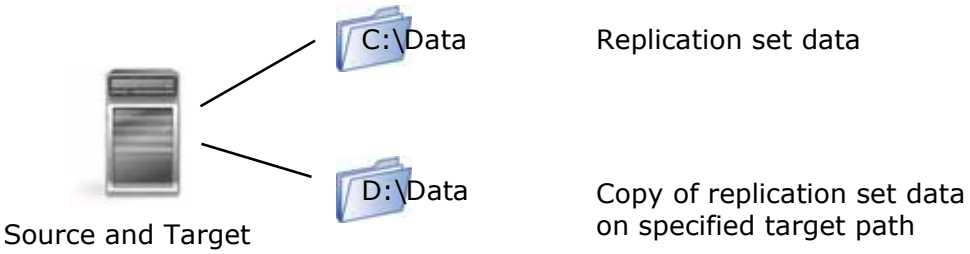
Description	One source server sends data to multiple target servers. The target servers may or may not be accessible by one another.
Applications	This configuration provides offsite disaster recovery, redundant backups, and data distribution. For example, this configuration can replicate all data to 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.

Chained



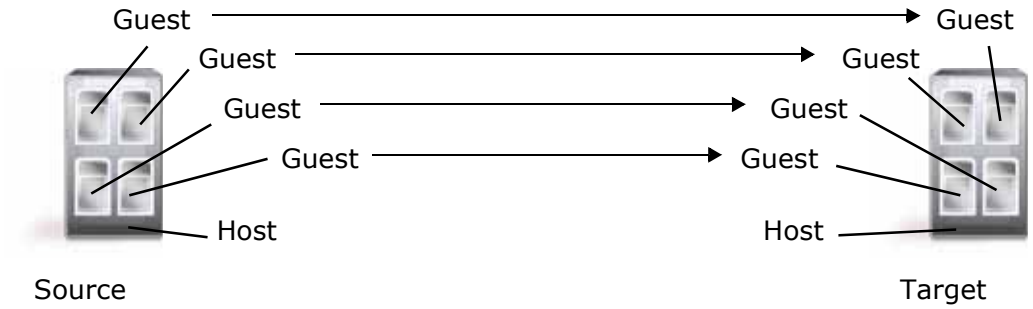
Description	The source servers sends replicated data to a target server, which acts as a source server and sends data to a final target server, which is often offsite.
Applications	This is a convenient approach for integrating local high availability with offsite disaster recovery. This configuration 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.

Single server

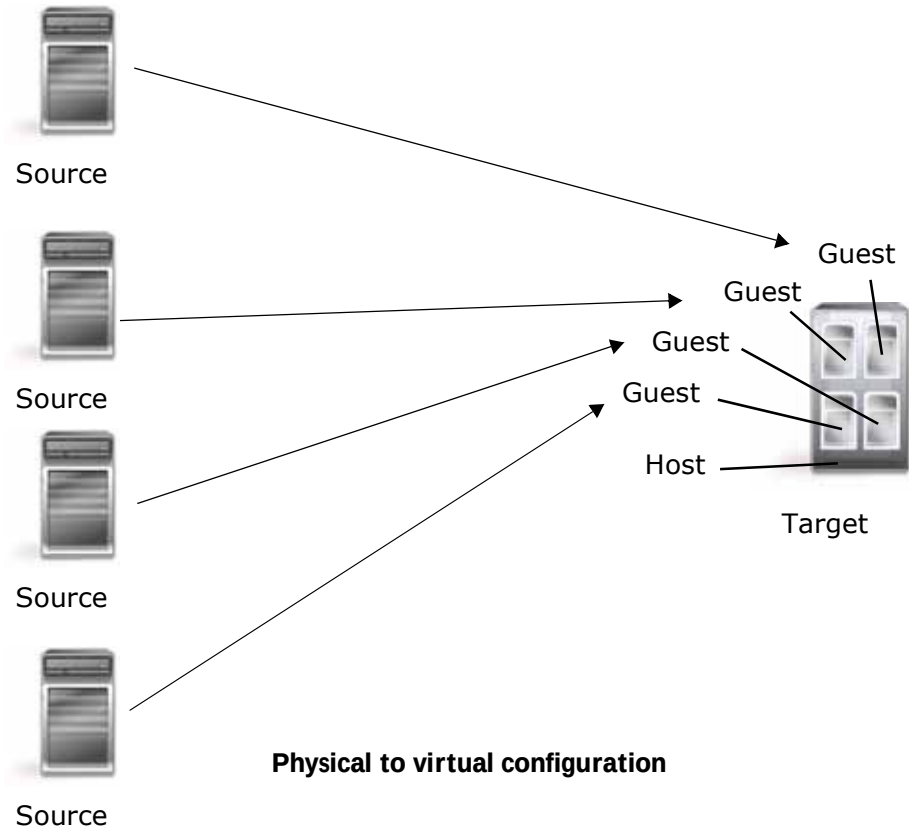


Description	Source and target Double-Take components are loaded on the same server allowing data to be replicated from one location to another on the same volume or to a separate volume on the same server. These could be locally attached SCSI drives or Fibre Channel based SAN devices.
Applications	This configuration is useful for relocating data within existing servers so that open files can be backed up. If the source and target copies of the data are located on different drives, this configuration does support high availability of the replication set in the event that the source hard drive fails.
Considerations	This configuration does not provide high availability for the entire server. This configuration must be configured carefully so that an infinite loop is not created.

Virtual systems



Virtual to virtual configuration



Physical to virtual configuration

Description	Source and/or target Double-Take components are loaded on the guest operating system allowing data to be replicated between virtual systems. One Double-Take for Virtual Systems® license can be used on up to five virtual servers.
Applications	This configuration easily protects virtual environments. When combined with physical servers, this configuration offers additional flexbility.
Considerations	This configuration does not protect data outside of the virtual system.

Resources

You have many resources available to you when using Double-Take.

- **Operating System and application documentation**—Make sure that you have complete documentation for your operating system and your applications.
- **Double-Take documentation**—The complete set of Double-Take documentation includes the manuals listed below. Each manual is available in the \DblTake\docs directory on the product CD and in the root of the installation directory you selected during the installation. The manuals are in Adobe® Acrobat® PDF format. If needed, you can install the free Adobe Acrobat Reader® by downloading the latest version from the Adobe web site at www.adobe.com.
  - *Double-Take for Linux User's Guide*—This guide is a complete how-to guide with detailed steps for installing and using Double-Take. This guide also contains a complete list of Double-Take messages, alerts, and statistics. The file name of this manual is User's Guide.pdf.
  - *Double-Take for Linux Scripting Guide*—This guide is a scripting manual for those users who want to use DTCL (Double-Take Command Language). The file name of this manual is Scripting Guide.pdf.
  - *Double-Take for Linux Evaluation Guide*—This guide is for those users who are evaluating Double-Take. It walks through specific steps allowing you to become familiar with the Double-Take core technology first hand. The file name of this manual is Evaluation Guide.pdf.
- **Double-Take Online Help**—The Double-Take Management Console and Failover Control Center contain context-sensitive online help. The online help can be accessed by pressing the F1 key, clicking the **Help** button on screens where it is available, or selecting **Help**, **Help Topics**.
- **Readme**—Double-Take includes a readme file which contains any last-minute information. This file is called readme.html and is located in the \DblTake\docs directory of the product CD and in the root of the installation directory you selected during the installation.
- **Application Notes**—Double-Take Software tests many of the popular applications on the market today and the results of these testing procedures are written up into formal Application Notes that describe how Double-Take should be configured to work correctly with certain applications. For a complete list of Application Notes, see the Double-Take Software Support web site at support.doubletake.com.
- **Other Resources**—Technical support, professional services, and training are all available for all Double-Take Software products.
  - **Technical Support**—If you need technical assistance, you can contact technical support. All basic configurations outlined in the documentation will be supported through technical support. Your technical support center is dependent on the reseller or distributor you purchased your product from and is identified on your service agreement. If you do not have access to this agreement, contact Double-Take Software Technical Support and they will direct you to the correct service provider. To contact Double-Take Software Technical Support, you will need your serial number and activation code. For a complete list of US and international phone numbers, see the Double-Take Software web site at www.doubletake.com. Online support is available at support.doubletake.com.
  - **Professional Services**—Assistance and support for advanced configurations may be referred to a Pre-Sales Systems Engineer or to Professional Services. See the Double-Take Software web site at www.doubletake.com for more information.
  - **Training**—Classroom and computer-based training are available. See the Double-Take Software web site at www.doubletake.com for more information.

Installation

2

Review ***System requirements*** and then use the appropriate instructions from the following list to meet your goal.

- ***Installing or upgrading Double-Take for Linux*** on page 2-3—Use these instructions if you are installing or upgrading on a Linux server.
- ***Configuring your servers using DTSetup*** on page 2-4—Use these instructions to configure your Double-Take servers.
- ***Installing or Upgrading the Double-Take Client for Windows*** on page 2-9—Use these instructions if you want to install or upgrade Double-Take on a Windows® client.
- ***Removing (Erasing) Double-Take*** on page 2-10—Use these instructions if you want to remove Double-Take from a Linux server.
- ***Removing the Double-Take Client for Windows*** on page 2-10—Use these instructions if you want to uninstall Double-Take from a Windows client.

System requirements

Verify that each machine that will be used as a Double-Take source or target meets the following system requirements.

- **Operating System**—Make sure your servers meets the operating system requirements outlined in the following table.

Physical Servers and Virtual Systems Guest Operating Systems	Kernel Version	Kernel Type	File System	Virtual Systems Platforms
• Red Hat® Enterprise Linux 4.5 through 4.7 for 32-bit architecture • CentOS 4 Updates 5-7 for 32-bit architecture	2.6.9.x.x	• Default • SMP • HugeMem	• Ext2 • Ext3	• VMware® • Microsoft Virtual Server • Xen® • Virtual Iron®
• Red Hat Enterprise Linux 4.5 through 4.7 for 64-bit architecture • CentOS 4 Updates 5-7 for 64-bit architecture		• Default • SMP • LargeSMP		
• Red Hat Enterprise Linux 5.0 through 5.3 for 32-bit architecture • CentOS 5 GA-Update 2 for 32-bit architecture	2.6.18.x.x	• Default (SMP) • Xen • PAE		
• Red Hat Enterprise Linux 5.0 through 5.3 for 64-bit architecture • CentOS 5 GA-Update 2 for 64-bit architecture		• Default (SMP) • Xen		

Physical Servers and Virtual Systems Guest Operating Systems	Kernel Version	Kernel Type	File System	Virtual Systems Platforms
SUSE Linux Enterprise Server 10.2 for 32-bit architecture	2.6.16.x.x	• Default • SMP • BigSMP • Xen • XenPAE	• Ext2 • Ext3 • ReiserFS	• VMware® • Microsoft Virtual Server • Xen® • Virtual Iron®
SUSE Linux Enterprise Server 10.2 for 64-bit architecture		• Default • SMP • Xen		

NOTE: Because Double-Take 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 installed, you must remove Double-Take prior to the operating system upgrade. Uninstall Double-Take, perform the operating system upgrade, and then reinstall Double-Take.

- **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 program files is approximately 45 MB. About 5 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 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**—TCP/IP
- **Name resolution**—Your servers must have name resolution using hosts files entries, DNS, and/or DHCP. For details on configuring name resolution by any of these methods, see your Linux documentation or online Linux resources.

Installing or upgrading Double-Take for Linux

Use these instructions if you are installing or upgrading Double-Take.

1. Determine the installation package that is appropriate for your operating system.

Operating System	Installation File
32-Bit Operating System	DoubleTake-versionnumbers.i386.rpm
64-Bit Operating System	DoubleTake-versionnumbers.x86_64.rpm

   The versionnumbers in the installation file name will vary and will correspond to the version of Double-Take you are installing. For example, if you are installing version 4.5.11.0.2200, the installation files would be DoubleTake-4.5.11.0-2200.i386.rpm or DoubleTake-4.5.11.0-2200.x86_64.rpm. If you are uncertain about the architecture of your machine, you can use the uname -m command to determine it.

2. 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. Go to a shell prompt by launching a terminal application from your UI or logging in via the Linux virtual console.
     b. If you are not already, become the superuser ("root") by typing su -.
     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-4.5.11.0-2200.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.

3. After you have installed the software, you need to configure your installation prior to using Double-Take. Run the following command from the shell prompt.

   DTSetup

NOTE: 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.

4. The first time you run DTSetup after an installation or upgrade, you will be prompted to review the Double-Take Software license agreement. Review the agreement and accept the terms of agreement by typing yes. You cannot use Double-Take without agreeing to the licensing terms.
5. When the DTSetup menu appears, at a minimum you must specify an activation code, configure replication, and start the daemon. You may also desire to modify the security group configuration. During the install, root is automatically added to the main Double-Take security group. If you want to add other users or remove root, you should modify the security group configuration as well. For details on all of the DTSetup menu options, see ***Configuring your servers using DTSetup*** on page 2-4.

Configuring your servers using DTSetup

DTSetup is a menu-driven application to provide easy access to Double-Take server configuration.

1. Run the following command from the shell prompt to start DTSetup.

   DTSetup

NOTE:	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 or upgrade, you will be prompted to review the Double-Take Software license agreement. Review the agreement and accept the terms of agreement by typing yes. You cannot use Double-Take without agreeing to the licensing terms.
3. When the DTSetup menu appears, enter the number of the menu option you want to access. Use the following list (with corresponding menu numbers) to understand each menu option.

```
=== DTSetup Main Menu ===


Menu Options:
1.    Setup tasks
2.    Start/Stop Double-Take daemon
3.    Start User Interface (DTCL -i)
4.    Documentation/Troubleshooting tasks
Q.    Quit DTSetup


Please choose a menu option :
```

1. **Setup tasks**—Several setup tasks are required for Double-Take functionality, including an activation code for each server, at least one user in the Double-Take administrator group, and replication configuration for your source. The remaining items in the setup menu are also Double-Take configuration related, but are optional.
   1.1. **Set activation code**—For each server, you must specify a unique activation code, which is an alpha-numeric code that applies the appropriate Double-Take license to your installation.
   1.2. **Add/Remove users to Double-Take groups**—Double-Take uses the native operating system to provide security. Privileges are granted through membership in user groups defined on each machine running Double-Take. To gain access to a particular Double-Take 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 security groups. The administrator group (dtadmin) provides full privileges and the monitor group (dtmon) provides a subset of privileges. During the install, root is automatically added to the Double-Take administrator group. See ***Security*** on page 17-1 for details on security and the specific privileges granted to each group.
      1.2.1. **Add Double-Take Administrator**—This option adds users to the Double-Take administrator group.
      1.2.2. **Remove Double-Take Administrator**—This option removes users from the Double-Take administrator group
      1.2.3. **Add Double-Take Monitor**—This option adds users to the Double-Take monitor group.
      1.2.4. **Remove Double-Take Monitor**—This option removes users from the Double-Take monitor group.

1.3. **Configure File System or Block Device Replication**—Double-Take hooks operations for replication by interjecting itself between the kernel VFS layer and the underlying file system. This must be done on a case by case basis for each block device, mount point, or parent directory that needs to be replicated.

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

1.3.1. **Configure File System Replication setup file**—In order to use Double-Take replication, data that will be replicated must reside on a partition, which is specially mounted with the Double-Take file system (DTFS) driver. DTFS is a transparent file system that monitors the storage file system so that data changes can be captured. The data on those mount points can only be replicated if they are under DTFS mount points. The DTSetup utility allows you to configure entries in /etc/DT/dtfs_mounts to mount specific partitions as DTFS when Double-Take is started.

NOTE: Double-Take is compatible with NFS and Samba services as long as they are mounted on top of Double-Take. Double-Take should be mounted just above the storage file system. Additionally, NFS and/or Samba must 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.

1.3.1.1. **Add an entry to /etc/DT/dtfs_mounts**—Type 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. You can either stop and restart the Double-Take daemon so that new entries in /etc/DT/dtfs_mounts will take effect or you can mount the entries immediately using menu option 1.3.1.3.

NOTE: 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.)

1.3.1.2. **Remove an entry from /etc/DT/dtfs_mounts**—Type the path of the entry that needs to be removed.

1.3.1.3. **Mount all entries in /etc/DT/dtfs_mounts as DTFS**—This option will cause the entries of the list to be mounted with DTFS immediately, rather than when the daemon is started. This allows you to review the effects of the changes and troubleshoot any issues.

1.3.1.4. **Unmount DTFS from all entries in /etc/DT/dtfs_mounts**—This will unmount DTFS from the mount points on the system, if any. This allows you to make changes to the DTFS mounts without unloading the daemon.

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

1.3.2. **Configure Block Device Replication setup file**—In order to use Double-Take 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 loop driver (DTLOOP). DTLOOP allows the loop device to serve as a monitoring 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 system.

1.3.2.1. **Add an entry to /etc/DT/dtloop_devices**—Enter the path to the block device that is to be replicated and press **Enter**. Then, enter the path to the loop device to use (/dev/loop#), if the same one should always be attached. 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. After pressing **Enter**, 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 documentation for losetup for more information on using encryption. You can either stop and restart the Double-Take daemon so that new entries in /etc/DT/dtloop_devices will take effect or you can mount the entries immediately using menu option 1.3.2.3.

1.3.2.2. **Remove an entry from /etc/DT/dtloop_devices**—Enter the path of the block device that needs to be removed.

1.3.2.3. **Attach all entries in /etc/DT/dtloop_devices to a loop device**—This option will cause the entries of the list to be attached to loop devices with DTFS immediately, rather than when the daemon is started. This allows you to review the effects of the changes and troubleshoot any issues.

1.3.2.4. **Detach loop devices from all entries in /etc/DT/dtloop_devices**—This option will detach the DTLOOP loop device from the block devices on the system, if any. This allows you to make changes to the DTLOOP loop devices without unloading the daemon.

1.3.2.5. **List block devices on this system**—This option will provide a list of block devices to which DTLOOP can be attached. Use this list as a guide for menu option 1.3.2.1.

1.3.3. **Manual Replication Configuration menu**—This option will take you to the menu where the file system and block device replication configuration can be changed but not persisted between reboots/restarts.

1.3.3.1. **Manually mount DTFS on a file system**—Type 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.

NOTE:	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.)

1.3.3.2. **Manually unmount a DTFS file system**—Type the path of the entry that needs to be removed from replication.

1.3.3.3. **Manually attach DTLOOP to a block device**—Enter the path to the block device that is to be replicated and press **Enter**. Then, enter the path to the loop device to use (/dev/loop#). DTLOOP can also use the first one available. After pressing **Enter**, 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 documentation for losetup for more information on using encryption.

1.3.3.4. **Manually detach DTLOOP from a block device**—Enter the path of the block device that needs to be removed from replication.

1.4. **Configure iptables for Double-Take**—If you are using Double-Take across a router, you will need to configure the router ports to accommodate Double-Take traffic. Use this option to assist in IP address configuration.

1.4.1. **Open Double-Take ports in iptables**—This option attempts to open the IP ports needed by Double-Take: 1500 (TCP/UDP), 1505 (UDP), and 1506 (TCP). This option only works on RedHat firewalls with a default configuration and is provided for administrators who are not familiar with configuring firewalls. If possible, you should open the port yourself.

1.4.2. **Close Double-Take ports in iptables**—This option attempts to close the ports that were opened in menu option 1.4.1.

1.5. **Edit Double-Take config file**—Double-Take program settings are the user-definable options that control Double-Take application settings and Double-Take processing settings. This includes options like the name of the verification log file and the size of the Double-Take memory pool. You can set these options using DTSetup or the Double-Take clients.

1.6. **Configure Double-Take driver performance**—You can configure various throttling and memory settings from DTSetup, although changes to these settings should be made with caution. Contact technical support for advice on the proper settings for your specific environment.

2. **Start/Stop Double-Take daemon**—The Double-Take daemon will start automatically after Double-Take is installed and the server is rebooted. You can start and stop the Double-Take daemon using this built-in DTSetup script.

2.1. **Start Double-Take and process driver config**—This option starts the Double-Take daemon and the auto replication configuration is processed.

2.2. **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 and unmount replicated devices.

2.3. **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 and unmount replicated devices.

2.4. **Restart service and reset driver config**—This option does a full stop and start, completely unloading the daemon and drivers and then reloading them. Do not unload the driver if there are still processes using the DTFS mount point.

2.5. **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. Do not unload the driver if there are still processes using the DTFS mount point.

2.6. **Go to Replication Configuration menu**—This option takes you to menu option 1.3.3 and then returns you to this point when exit that menu.

3. **Start User Interface** (DTCL -i)—You can launch the Double-Take interactive command prompt which allows you to enter DTCL commands one at a time.

4. **Documentation/Troubleshooting tasks**—This menu provides access to Double-Take log files, a diagnostic collection tool, user documentation, and several legal documents.

4.1. **View log files**—Double-Take logs alerts, which are proccessing notifications, warnings, and error messages. The logs are written to disk.

4.1.1. **View /var/log/DT/dtlog*.dtl in less**—This option uses the less file viewer program to view all of the Double-Take logs, starting from the most recent.

4.1.2. **Follow the output of latest dtlog**—This option uses tail -f to watch the output of the Double-Take logs in real-time.

4.1.3. **View /var/log/messages in less**—This option uses the less file viewer program to view the system log messages.

4.1.4. **Follow the output of /var/log/messages**—This option uses tail -f to watch the output of the system log messages in real-time.

4.2. **Collect and package diagnostic info**—You may need to collect configuration data when reporting problems to technical support. DTSetup can gather Double-Take log files, Double-Take settings, network configuration information such as IP addresses, and other data which may be necessary for technical support to troubleshoot issues.

4.3. **View user documentation**—You can view Double-Take documentation such as the Evaluation Guide, User's Guide, readme, and several legal documents. DTSetup attempts to determine your viewers.

4.3.1. **View ReadMe HTML**—This option views the readme file which contains last minute release notes.

4.3.2. **View Scripting Guide PDF**—This option views the *Scripting Guide* which contains DTCL commands and scripting information.

4.3.3. **View Users Guide PDF**—This option views the *User's Guide* which contains instructions for using Double-Take.

4.3.4. **View Evaluation Guide PDF**—This option views the *Evaluation Guide* which contains step-by-step instructions for evaluating Double-Take.

4.3.5. **View Premium Maintenance Agreement PDF**—This option views the Premium Maintenance Agreement legal document.

4.3.6. **View Enterprise Maintenance agreement PDF**—This option views the Enterprise Maintenance Agreement legal document.

4.3.7. **View End User License Agreement TXT**—This option views the End User License Agreement legal document.

4.3.8. **View dtfs/dtrep module license TXT**—This option views the open source legal document.

4.3.9. **Change a document viewer**—This option allows you to specify a document viewer.

Installing or Upgrading the Double-Take Client for Windows

Use the instructions to install the Double-Take clients on a Microsoft Windows machine. Your Windows version can be Windows 2003, Windows XP Service Pack 2 or later, or Windows Vista on 32-bit or 64-bit platforms.

1. Close any open applications.
2. Start the installation program using the appropriate instructions, depending on your media source.
   - **CD**—Load the Double-Take CD into the local CD-ROM drive. If auto-run is enabled, the installation program will start automatically. To manually start the program, select Start, Run and specify <cd_drive>:\autorun.exe. Select to install Double-Take for Windows.
   - **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 installation program begins, the Welcome screen is displayed. Click **Next** to continue.
4. Review and accept the Double-Take 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 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

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

1. Go to a shell prompt by launching a terminal application from your graphical user interface (GUI) or logging in via the Linux virtual console.
2. If you are not already, become the superuser ("root") by typing su -.
3. Erase Double-Take by using the command rpm - e DoubleTake.

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.

NOTE: After removing Double-Take, some files may not be removed until after the server has been rebooted

Removing the Double-Take Client for Windows

Use these instructions if you want to remove the Double-Take clients on a Microsoft Windows machine.

NOTE: If desired, you can use the Add/Remove Programs function of the Windows Control Panel to remove the Double-Take clients.

1. Close any open applications.
2. Start the installation program using the appropriate instructions, depending on your media source.
   - **CD**—Load the Double-Take CD into the local CD-ROM drive. If auto-run is enabled, the installation program will start automatically. To manually start the program, select Start, Run and specify <cd_drive>:\autorun.exe. Select to install Double-Take for Windows.
   - **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 installation program begins, the Welcome screen is displayed. Click **Next** to continue.
4. On the Maintenance dialog box, select **Remove** and click **Next** to continue.
5. The installation program is immediately ready to begin removing the Double-Take files. Click **Remove** to continue.

When the message prompt informs you that the uninstall program has completed successfully, click **OK**.

Clients

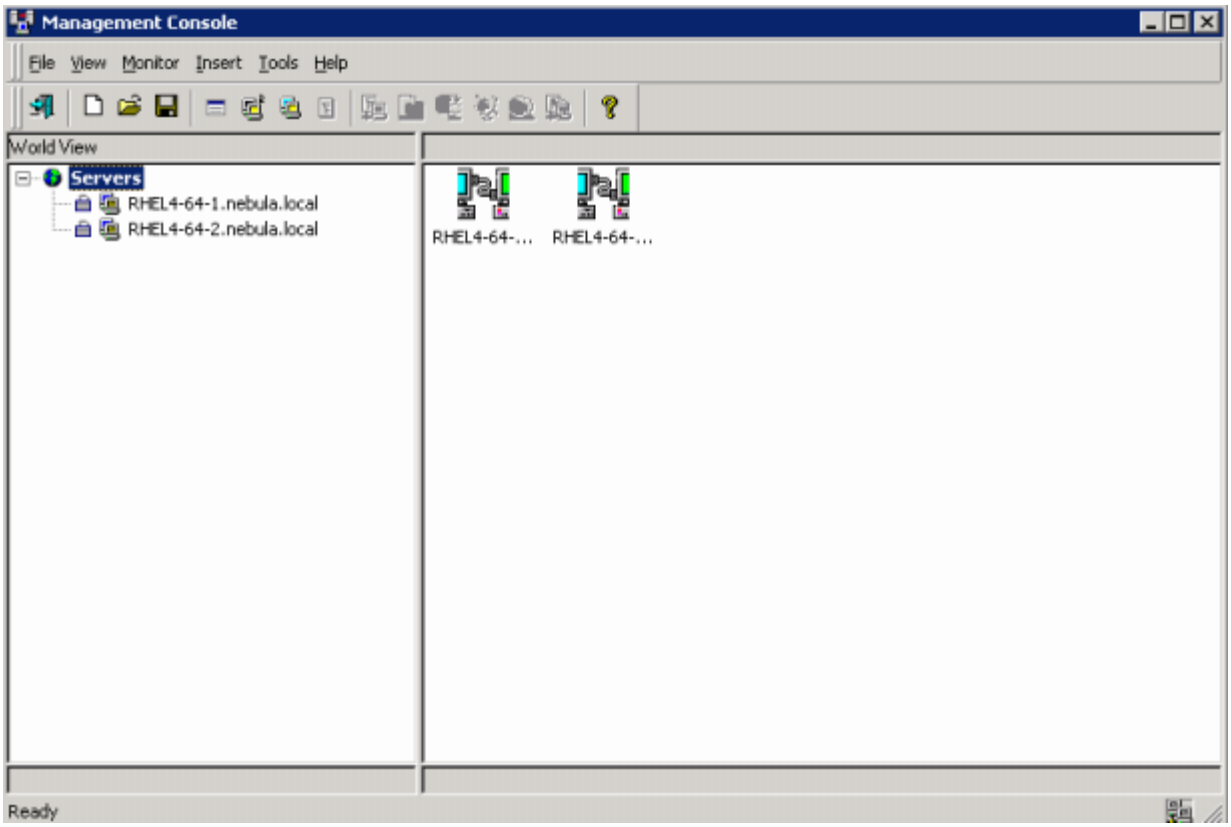
3

Double-Take has two clients, the Management 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 Management Console.

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

Management Console

From the Management Console, you can manage, monitor, and control your Double-Take connections. The Management 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 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.



Logging on and off of Double-Take

To ensure protection of your data, Double-Take offers multi-level security using native operating system security features. Privileges are granted through membership in user groups defined on each machine running Double-Take. To gain access to a particular Double-Take 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 security groups. Once a valid user name and password has been provided and the Double-Take source or target has verified membership in one of the Double-Take 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 is granted on one of the following three levels:

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

1. Highlight a machine on the left pane of the Management Console. By double-clicking the machine name, Double-Take 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. Click **OK**. Verify your access by the resulting icon and log on again if necessary.

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

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

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

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

Icon	Description	Access Granted
	This icon is a computer with a gear and it indicates the Double-Take security is set to administrator access.	Administrator rights
	This icon is a computer with a magnifying glass and it indicates the Double-Take security is set to monitor only access.	Monitor rights
	This icon is a lock and it indicates the Double-Take security is set to no access.	No rights

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

Double-Take workspaces

The Management Console workspace contains the display of the panes of the Management 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 Management Console, you can save the workspace to use later or use on another Double-Take 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.

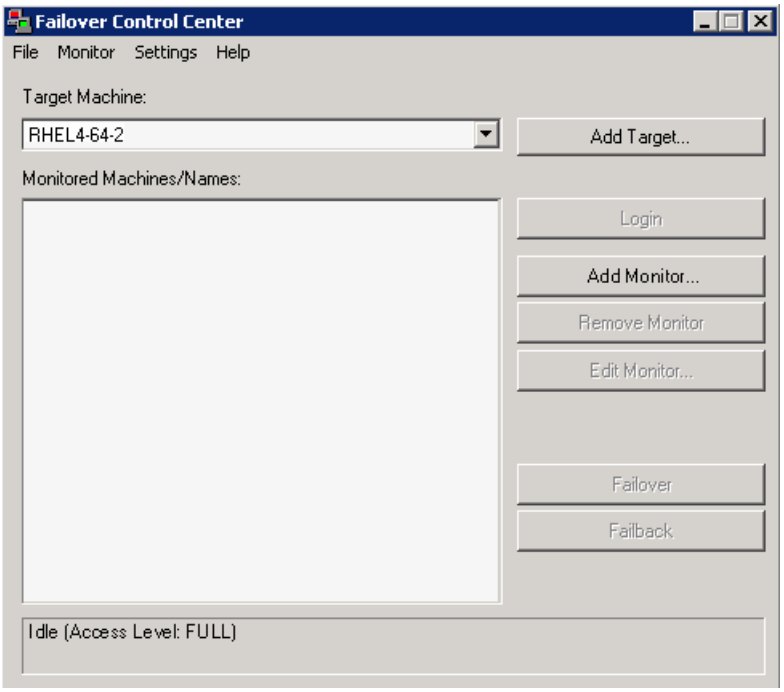
Opening a workspace

From the Management 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 window settings.
- **Open Workspace**—Open a previously saved workspace.

Failover Control Center

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



Configuring communication ports

The Failover Control Center use two ports for various Double-Take communications.

- **Service Transmit Port**—The Failover Control Center uses this port to send commands to Double-Take servers.
- **Heartbeat Listen Port**—The Failover Control Center uses this port to listen for heartbeats transmitted from Double-Take servers.

To view or modify the port settings in the Failover Control Center, select **Settings**, **Communications**.

Configuring refresh rate

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

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

Clearing maintained security credentials

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

Replication Sets

4

A replication set defines the data on a source machine that Double-Take 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 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 can mirror, replicate, verify, and restore paths up to 4,095 characters. Paths longer than 4,095 characters will be skipped and logged to the Double-Take 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 Management 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 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 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. 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 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.

Replication capabilities

Double-Take replicates all file and directory data in the supported Linux file systems. (See ***System requirements*** on page 2-1 for the file systems supported). It 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 partition, which is specially mounted with the Double-Take 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 point 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. See ***Configuring your servers using DTSetup*** on page 2-4 for more information on the DTSetup utility.
- DTFS mount points are identified in the Management Console by dtfs in the mount point name.
- Double-Take is compatible with NFS and Samba services as long as they are mounted on top of Double-Take. Double-Take 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 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 will remap the path contained in that link based on the Double-Take 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 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 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. Opertions on hardl inks 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.

Creating a replication set

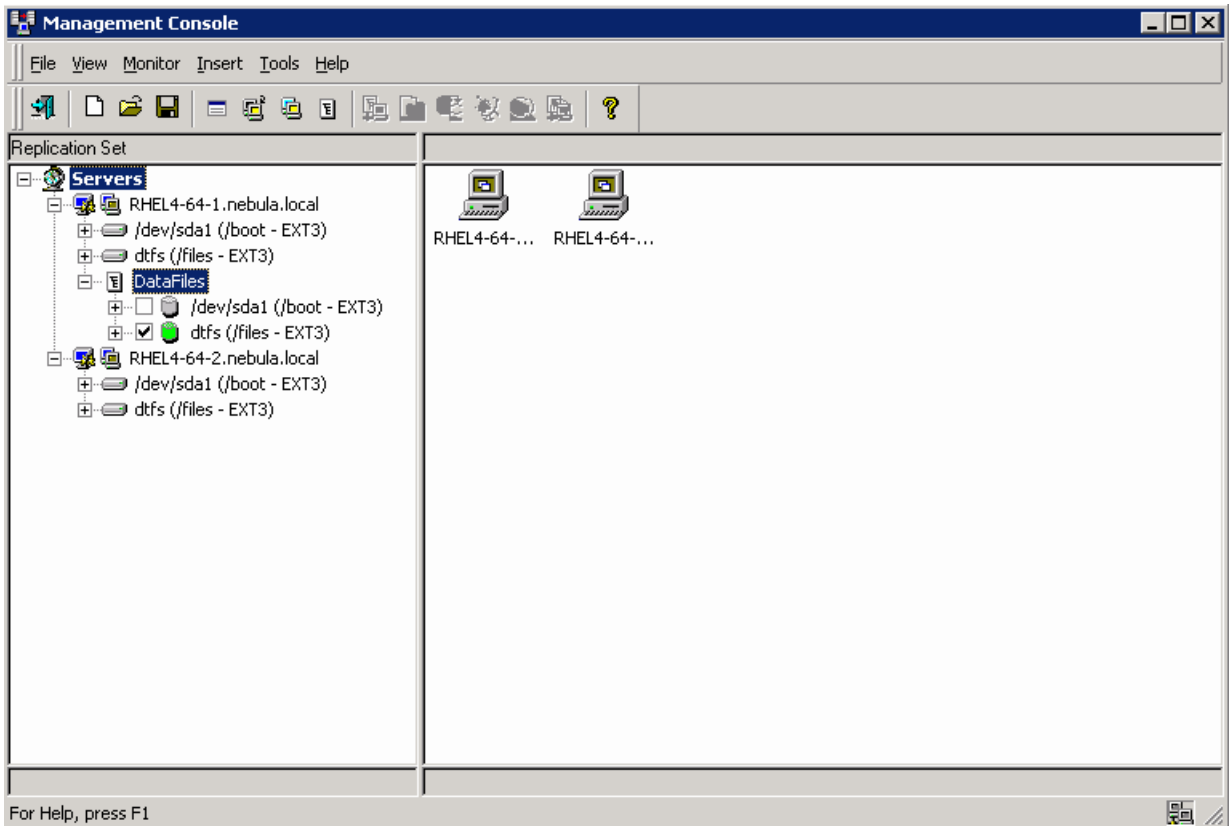
Before you can establish a connection, you must create a replication set.

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

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

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

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



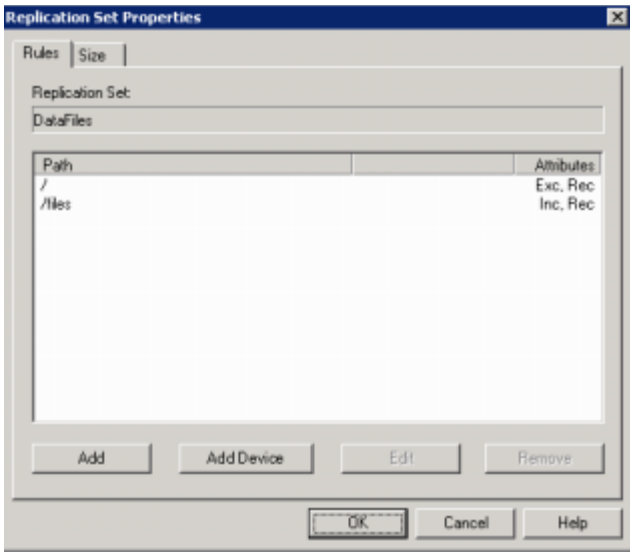
NOTE: Be sure and verify what files can be included by reviewing ***Replication capabilities*** on page 4-3.

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

Creating or modifying replication rules manually

There may be times when you cannot browse for data when creating a replication set. For example, you can create a replication set rule for a directory or file that does not exist. Since you cannot browse for the location, you have to create replication set rule manually. At other times, the data you want to replicate cannot be easily selected from the Management 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 Management 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 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 replication. You can select any of these devices, but you cannot use them for Double-Take replication until they are configured for Double-Take replication. The status no dtloop indicates the device is not configured for Double-Take replication. See ***Configuring your servers using DTSetup*** on page 2-4 for details on replication configuration.
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 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 replication are shown by default. Highlight the device to include in the replication set and click **OK**.

NOTE:	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 replication. You can select any of these devices, but you cannot use them for Double-Take replication until they are configured for Double-Take replication. The status **no dtloop** indicates the device is not configured for Double-Take replication. See ***Configuring your servers using DTSetup*** on page 2-4 for details on replication configuration.

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

Modifying a replication set

Double-Take 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.

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

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.

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

Deleting a replication set

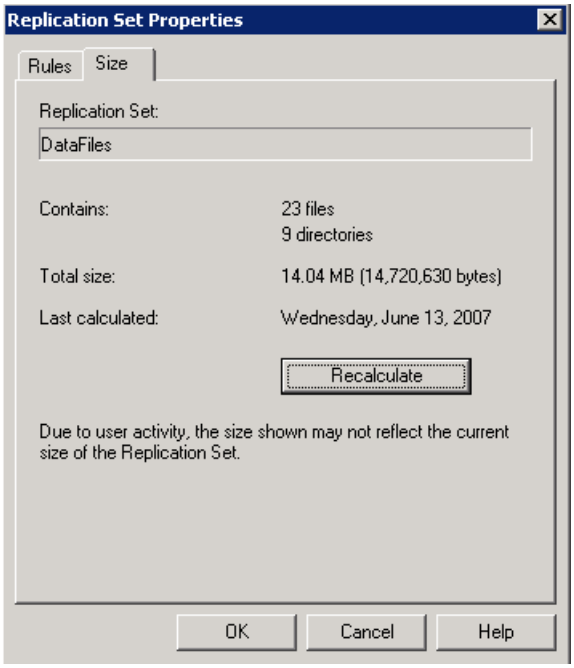
You can only delete a replication set if it is not currently connected. If the replication set is connected, you must disconnect the connection and then delete the replication set.

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.

Calculating replication set size

While Double-Take is mirroring, the right pane of the Management 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 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 Management Console.

NOTE: You can also configure the replication set calculation when establishing a connection through the Connection Manager by selecting **Calculate Replication Set size on connection** on the **Mirroring** tab.

If your replication set contains a large number of files, for example, ten thousand or more, you may want to disable the calculation of the replication set size so that data will start being mirrored sooner. If calculation is enabled, the source calculates the file size before it starts mirroring. This can take a significant amount of time depending on the number of files and system performance. Disabling calculation will result in the mirror status not showing the percentage complete or the number of bytes remaining to be mirrored.

Establishing a Connection Using the Connection Wizard

5

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. The connection itself provides disaster recovery protection. If you want to extend the protection and provide high availability, you can add failover monitoring.

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

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

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

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

3. If you highlighted a source in the Management Console, the source will already be selected. If it is not, select the Double-Take source. This is the server that you want to protect. Click **Next** to continue.

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

4. If you highlighted a target in the Management Console, the target will already be selected. If it is not, select the Double-Take target. This is your backup server that will protect the source. Click **Next** to continue.

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

5. 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.

   Click **Next** to continue.
6. 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. Click **Next** to continue.

7. 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.

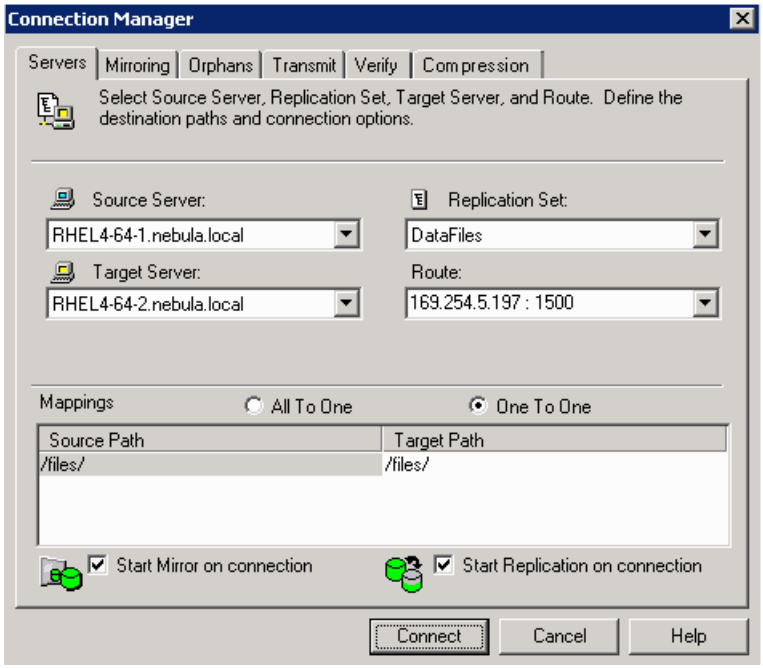
   Click **Next** to continue.
8. Review your selections on the summary screen. If your Connection Wizard settings are correct, establish your connection by completing one of the two options below:
   - If you do not want to set advanced options, click **Finish**. The Connection Wizard will close, the connection 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 Connection Manager will open. The **Servers** tab will be completed.

Establishing a Connection Using the Connection Manager

6

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 Management 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 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 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 traffic. For example, you can separate regular network traffic and Double-Take 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 offers two predefined locations as well as a custom option that allows you to create your own path.

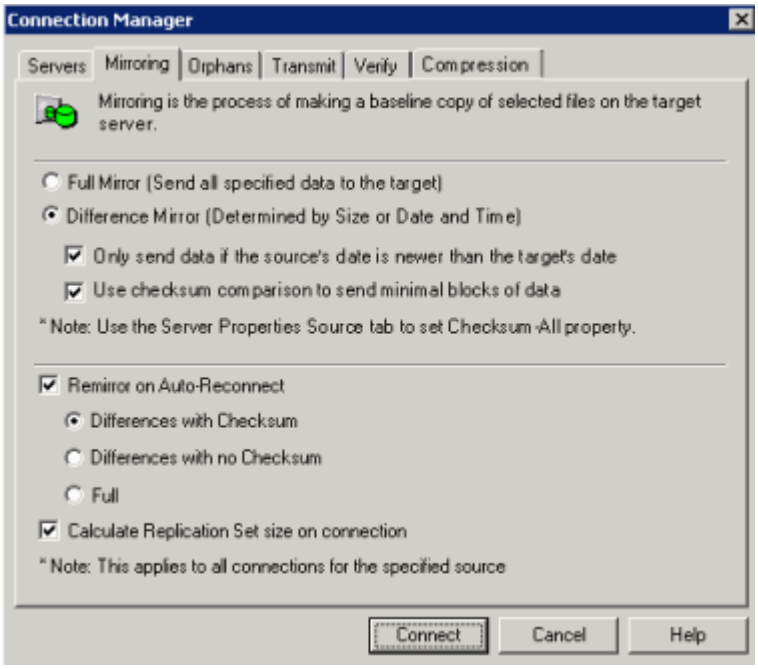
- **All To One**—This option replicates data from the source to a single volume on the target. The pre-defined path is /source_name/replication_set_name/volume_name. If you are replicating from multiple volumes on the source, each volume would be replicated to the same volume on the target.
- **One To 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.
- **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.

> **WARNING:** 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.

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

- **Use 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.

NOTE: See the table ***File Differences Mirror Options Compared*** on page 10-2 for 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 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 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.

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

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

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

Establishing a Connection Across a NAT or Firewall

7

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 coorelation to a single, external IP address. Double-Take 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.

Verifying the Double-Take port settings

Double-Take uses specific ports for communication between the Double-Take servers and Double-Take clients. In order to use Double-Take through a NAT or firewall, you must first verify the current Double-Take port settings so that you can open the correct ports on your hardware to allow Double-Take machines to communicate with each other.

Using the following table, locate and record your port settings for each of the Double-Take ports. The port setting can be found in the following locations.

- **Management Console**—From the Management Console, select **File**, **Options**, and go to the **Configuration** tab.
- **Failover Control Center**—From the Failover Control Center, select **Settings**, **Communications**.
- **Double-Take server**—From the Management Console, right-click on a server in the tree in the left pane of the Management Console, select **Properties**, and go to the **Network** tab.

Port	Description	Default Setting	Record Your Port Setting
Management Console Status Transmit Port	The Status Transmit Port sends and receives directed UDP comunications to display status and statistics in the Management Console.	1505	
Management Console Heartbeat Advertisement	The Heartbeat Advertisment port sends and receives broadcast UDP communications to populate the Management Console tree with Double-Take servers.	1500	
Failover Control Center Service Transmit Port	The Service Transmit Port sends and receives TCP communication between Double-Take servers and between Double-Take servers and Double-Take clients.	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 servers.	1500	
Double-Take Server Service Listen Port	The Service Listen Port sends and receives TCP communication between Double-Take servers and between Double-Take servers and Double-Take clients.	1500	

Port	Description	Default Setting	Record Your Port Setting
Double-Take Server Heartbeat Transmit Port	The Heartbeat Advertisement port sends and receives broadcast UDP communications to populate the Management Console tree with Double-Take servers.	1500	
Double-Take Server Status Listen Port	The Status Listen Port sends directed UDP comunications to display status and statistics in the Management Console.	1505	

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

Configuring the hardware

You need to configure your hardware so that Double-Take 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 server, its IP address, and the Double-Take and router ports. Also, note the following caveats.

- Since Double-Take communication occurs bidirectionally, make sure you configure your router for both incoming and outgoing traffic for all of your Double-Take servers and Double-Take clients.
- In addition to UDP heartbeats, Double-Take 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.

Inserting the servers in the Management Console

If your network is configured to propagate UDP broadcasts, your servers will be populated in the Management 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.

Establishing the connection

Once your server is inserted in the Management Console, you can use the Connection Wizard or the Connection Manager to establish your connection. See ***Establishing a Connection Using the Connection Wizard*** on page 5-1 or ***Establishing a Connection Using the Connection Manager*** on page 6-1 for details on establishing a connection through either of these methods.

Simulating a Connection

8

Double-Take 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.

Gathering system information

Before and after simulating your connection, you may want to gather network and system information specific to Double-Take operations. Use the DTSetup Diagnostics utility to automatically collect this data. It gathers Double-Take log files, Double-Take and system settings, network configuration information, and other data which may be necessary in evaluating Double-Take performance.

1. Run the DTSetup command from the shell prompt.

   NOTE: 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. Select option 4 for troubleshooting.
3. Select option 2 for diagnostics.
4. When you run DTSetup, it may take several mintues for it to finish processing. When it is complete, a .tar.gz file will be created in /var/run/DT. The file name will have DTInfo with the date and time.

Be sure and run the diagnostics before and after your simulation.

Establishing the simulated connection

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

1. 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**.
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 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.
3. Click **Connect** to establish the connection. The simulation data will be logged to the Double-Take statistics file. See ***DTStat*** on page 15-19 for more information.

Managing Connections 9

A unique connection ID is associated with each Double-Take 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 as described in the following table.

Connection Status	Description
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 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 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.

This chapter covers several connection related activities.

- ***Monitoring connections*** on page 9-2
- ***Monitoring the message window*** on page 9-6
- ***Queuing Double-Take data*** on page 9-8
- ***Configuring automatic reconnections*** on page 9-12
- ***Configuring compression*** on page 9-13
- ***Disconnecting a connection*** on page 9-13

Monitoring connections

When a source is highlighted in the left pane of the Management 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 servers, to help you monitor your connections. Review the *Connection Statistics* and *Connection and Server Display* tables for statistics and status descriptions. (Statistics marked with an asterisk (*) are not displayed, by default.)

Connection Statistics

Statistic	Description
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 daemon is restarted.
Target Name	The name of the target as it appears in the server tree in the left pane of the Management 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.

Statistic	Description
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 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.

Statistic	Description
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 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 Management 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.

You can change the statistics that are displayed by selecting **File**, **Options** and selecting the **Statistics** tab.

The statistics displayed in the Management 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.

The statistics appear on the Management 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.

Click **OK** to apply and save any changes that have been made to the order or display of the Management Console statistics.

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.

To return the statistics list to the Double-Take default selection and order, click **Reset to Default**.

Connection and Server Display

Location	Icon or Color	Description
Icons and colors 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	If the connection background is white, the Management Console and the source are communicating.
	Gray	If the connection background is gray, the Management 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.
Left pane icon		An icon with yellow and blue servers indicates a server that is working properly.
		A red X on a server icon indicates the Management 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.

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 Management 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 Management 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 Management 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 Management Console considers communications lost and updates the icons and colors. Valid values are between `0` and `1000`. The default setting is `5`.
5. Click **OK** to save the settings.

NOTE: If the **Site Monitor** and **Connection Monitor** settings are different, at times, the icons and color may not be synchronized between the left and right panes.

Monitoring the message window

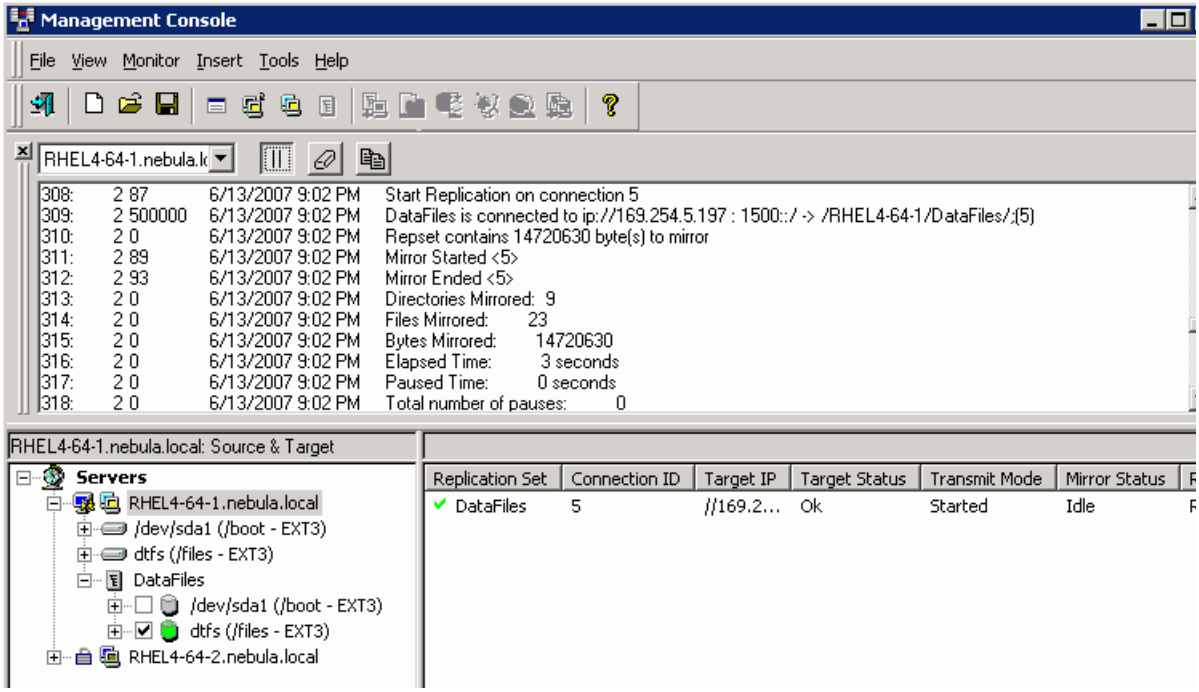
In addition to the statistics and status shown in the Management Console, you can also open a message window to view Double-Take processing alerts. These alerts include notifications, warnings, and errors.

NOTE: The standard appearance of the message window is a white background. If your message window has a gray background, the window is inactive. The Management 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 log on the server.

1. Open a new message window by 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.

   NOTE: Repeat this step to open multiple message windows.



2. To control the window after it is created, use either 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.

Message Window Control	Description	Toolbar Icon
Close	Closes the message window	
Clear	Clears the message window	

Message Window Control	Description	Toolbar Icon
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 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.	

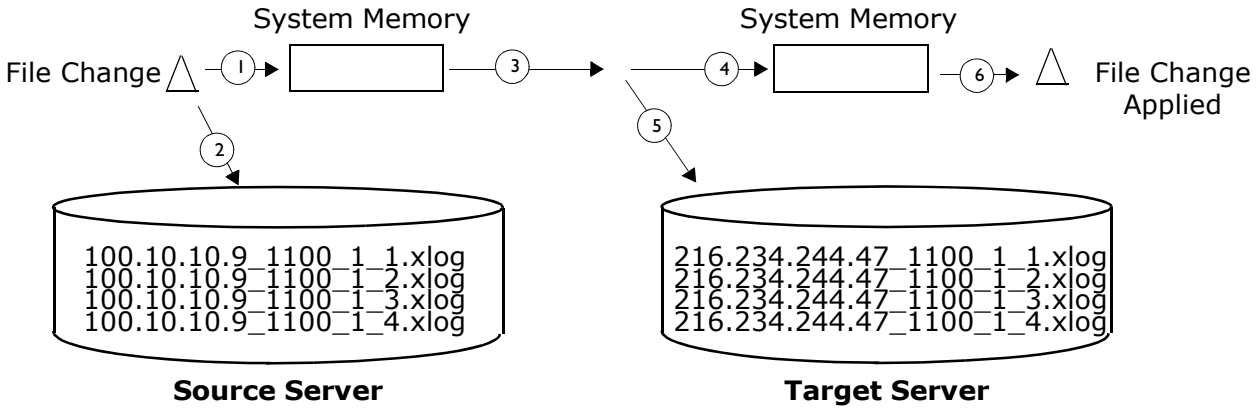
3. To change which server you are viewing messages for, select a different machine from the dropdown list on the toolbar. If necessary, the login process will be initiated.
4. To move the message window to other locations on your desktop, click and drag it to another area or double-click it to automatically undock it from the Management Console.

Queuing Double-Take data

During the Double-Take installation, you identified the amount of disk space that can be used for Double-Take queuing. Queuing to disk allows Double-Take 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.

Understanding queuing

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 port, the connection ID, and an incrementing sequence number.

   NOTE: You may notice transaction log files that are not the defined size limit. This is because data operations are not split. For example, if a transaction log has 10 KB left until the limit and the next operation to be applied to that file is greater than 10 KB, a new transaction log file will be created to store that next operation. Also, if one operation is larger than the defined size limit, the entire operation will be written to one transaction log.

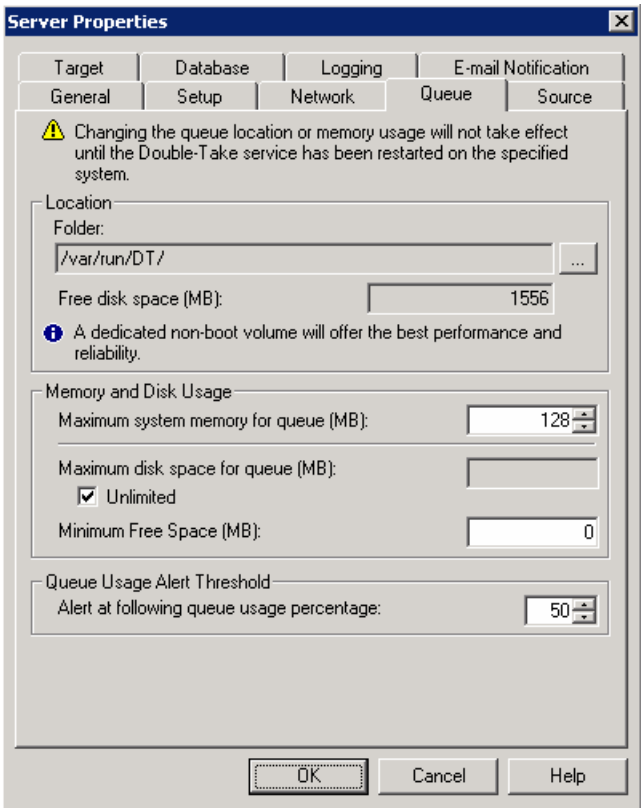
3. When system memory is full, the most recent changed data is added to the disk queue, as described in step 2. This means that system memory contains the oldest data. Therefore, when data is transmitted to the target, Double-Take pulls the data from system memory and sends it. This ensures that the data is transmitted to the target in the same order it was changed on the source. Double-Take 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 port, the connection ID, and an incrementing sequence number.

6. Like the source, system memory on the target contains the oldest data so when data is applied to the target, Double-Take pulls the data from system memory. Double-Take 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).

Configuring queuing

You should configure queuing on both the source and target.

1. Right-click the server on the left pane of the Management Console and select **Properties**.
2. Select the **Queue** tab.



3. Specify the **Folder** where the disk queue will be stored. Double-Take 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.

   NOTE: Scanning the Double-Take 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 queues.

4. Specify the **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 will use less system memory, but you will queue to disk sooner which may impact system performance. If you set it higher, Double-Take 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 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 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, you might configure the application to use 512 MB and Double-Take 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.

5. Specify the **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 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 will automatically begin the auto-disconnect process. By default, Double-Take will use an unlimited amount of disk space. Setting this value to zero (0) disables disk queuing.
6. Specify the **Minimum Free Space**. This is the minimum amount of disk space in the specified **Folder** that must be available at all times. By default, 50 MB of disk space will always remain free. The **Minimum Free Space** should be less than the amount of physical disk space minus **Maximum disk space for queue**.

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

7. 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%.
8. Click **OK** to save the settings.

Configuring automatic reconnections

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 will automatically start disconnecting connections. This is called auto-disconnect. The transaction logs and system memory are flushed allowing Double-Take to begin processing anew. The auto-reconnect process ensures that any connections that were auto-disconnected are automatically reconnected. Then, if configured, Double-Take 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 will auto-reconnect. If configured, Double-Take 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 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 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 source, not being able to communicate with the Double-Take 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 attemptto 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 is able to successfully persist this information, when the Double-Take daemon on the target is restarted, Double-Take will pick up where it left off, without requiring an auto-disconnect, auto-reconnect, or auto-remirror. If Double-Take 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 will auto-remirror. The remirror re-establishes the target baseline to ensure data integrity, so disabling auto-remirror is not advised.

NOTE: If you are experiencing frequent auto-disconnects, you may want to increase the amount of disk space on the volume where the Double-Take queue is located or move the disk queue to a larger volume. See ***Queuing Double-Take data*** on page 9-8 for more information.

If you have changed data on the target while not failed over, for example if you were testing data on the target, Double-Take is unaware of the target data changes. You must manually remirror your data from the source to the target, overwritting the target data changes that you caused, to ensure data integrity between your source and target.

Use the following steps to configure automatic reconnections.

1. Right-click the source server on the left pane of the Management 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.

Configuring compression

To help reduce the amount of bandwidth needed to transmit Double-Take 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 uncompresses it and then writes it to disk. On a default Double-Take installation, compression is disabled.

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

Keep in mind that the process of compressing data impacts processor usage on the source. If you notice an impact on performance while compression is enabled in your environment, either adjust to a lower level of compression, or leave compression disabled. Use the following guidelines to determine whether you should enable compression:

- If data is being queued on the source at any time, consider enabling compression.
- If the server CPU utilization is averaging over 85%, be cautious about enabling compression.
- The higher the level of compression, the higher the CPU utilization will be.
- Do not enable compression if most of the data is inherently compressed. Many image (`.jpg`, `.gif`) and media (`.wmv`, `.mp3`, `.mpg`) files, for example, are already compressed. Some images files, such as `.bmp` and `.tif`, are uncompressed, so enabling compression would be beneficial for those types.
- Compression may improve performance even in high-bandwidth environments.
- Do not enable compression in conjunction with a WAN Accelerator. Use one or the other to compress Double-Take data.

1. Right-click the connection on the right pane of the Management 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.

Disconnecting a connection

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

NOTE: If a connection is disconnected and the target is monitoring the source for failover, you will be prompted if you would like to continue monitoring for a failure. If you select **Yes**, the Double-Take 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 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 Management Console may become unresponsive while the data is being flushed. The Management Console will respond when all of the data has been flushed from the queue.

Mirroring and Replication 10

Mirroring and replication are the key components of Double-Take. This chapter contains information on the following mirroring and replication capabilities.

- ***Controlling mirroring*** on page 10-1
- ***Controlling an automatic mirror*** on page 10-3
- ***Configuring and removing orphan files*** on page 10-5
- ***Starting replication*** on page 10-6

Controlling 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 Management Console and select **Mirroring** and the appropriate mirror control.

- **Pause** or **Resume**—When pausing a mirror, Double-Take 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 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.
  - **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. See the table ***File Differences Mirror Options Compared*** on page 10-2 for 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**.
    - **Only send data only if source's date is newer than the target's date**—Only those files that are newer on the source are sent to the target.

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

    - **Use Checksum 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.
  - **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 Differences Mirror Options Compared

Server Properties Source Tab[a]	Connection Manager Mirroring Tab or Start Mirror Dialog			Action Performed
Checksum All	**File Differences**	**Only if Source is Newer**	**Checksum**	
(✓)	✓			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.
(✓)	✓	✓		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.
	✓		✓	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.
✓	✓		✓	The mirror performs a checksum comparison on all files and only sends those blocks that are different.
(✓)	✓	✓	✓	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.

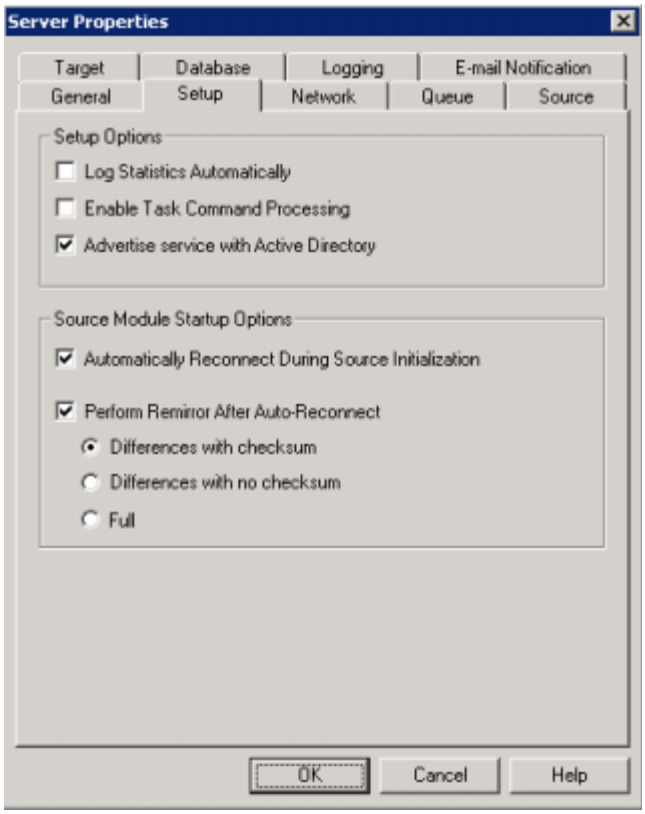
a. A checkmark enclosed in parentheses (✓) indicates that the global option on the Server Properties Source tab can be on or off. The use of this option does not change the action performed during the mirror.

Controlling an automatic mirror

In certain circumstances, for example if the disk-based queues on the source are exhausted, Double-Take 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 will perform an automatic mirror (called an auto-remirror) after an auto-reconnect.

NOTE: Auto-remirror is a per source option. When enabled, all connections from the source will perform an auto-remirror after an auto-reconnect. When disabled, none of the connections from the source will perform an auto-remirror after an auto-reconnect.

1. Right-click a server in the left pane of the Management 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.

NOTE: If auto-remirror is disabled and an auto-reconnect occurs, the transmission state of the connection will remain pending after the reconnect until a mirror is started manually.

4. 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.

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

See the table ***File Differences Mirror Options Compared*** on page 10-2 for 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.

Configuring and 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 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 log on the target, including the number of moved/deleted orphan files, the directories, and the number of bytes.

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

The orphans feature does not move or delete alternate data streams. To do this, use a full mirror which will delete the additional stream(s) when the file is re-created.

If Double-Take is configured to move orphan files, the Double-Take log file will indicate that orphan files have been deleted even though they have actually been moved. This is a reporting issue only.

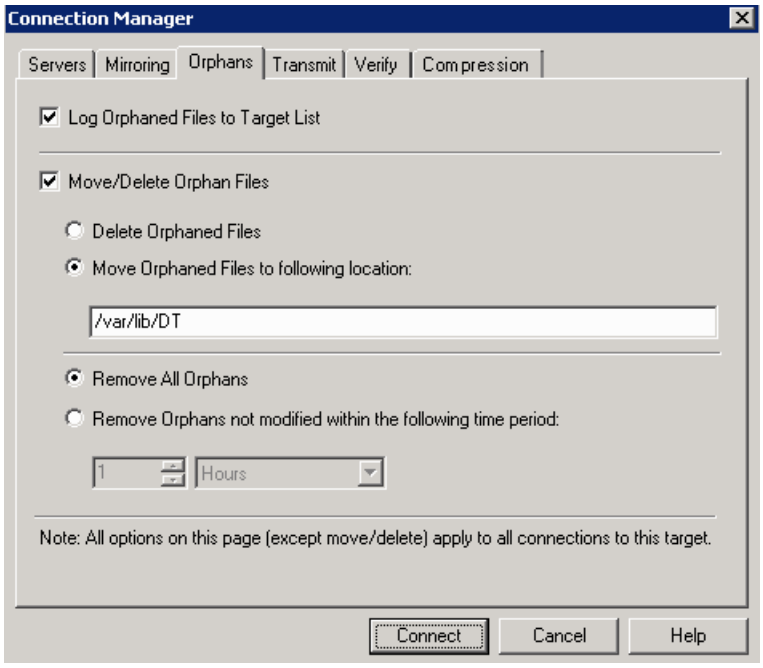
If delete orphans is enabled, directories and files that do not exist on the source and are excluded in the replication set using a wildcard rule will be removed from the target path. If you have data in your target path that does not exist on the source, do not use wildcard rules in your replication set. Manually select and deselect those files which should be included or excluded from your replication set.

If you want to remove orphan files manually, right-click an established connection and select **Remove Orphans**, **Start**. If you want to stop the process after it has been started, right-click the connection and select **Remove Orphans**, **Stop**.

To configure orphan files for processing during a mirror, verify, or restore, use the following instructions.

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

3. Specify if you want to log the name of the orphan files to the Double-Take log file on the target by marking **Log Orphaned Files to Target Log**.



4. By default, the orphan files feature is disabled. To enable it, mark **Move/Delete Orphan Files**.
5. Specify if you want to **Delete Orphaned Files** or **Move Orphaned Files to** a different location. If you select the move option, identify the location where these orphan files will be located.

NOTE:	If you are moving files, make sure the directory you specify to move the files to is not included in the destination of the replication set data so that the orphan files are only moved once.

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

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 Management Console and select **Replication**, **Start**. After starting replication, you should perform a remirror to guarantee the source and target data are identical.

Verification

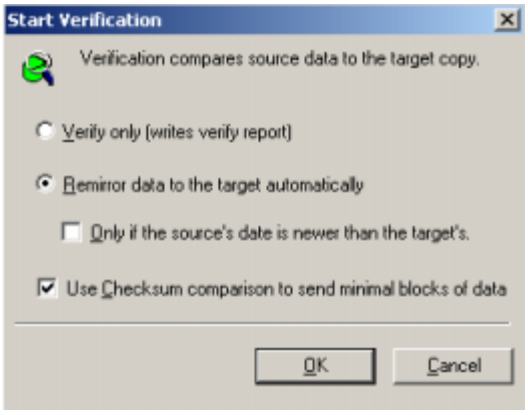
11

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 can automatically initiate a remirror. The remirror ensures data integrity between the source and target.

Manual verification

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

1. Right-click the connection on the right pane of the Management 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.

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

- **Use 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.

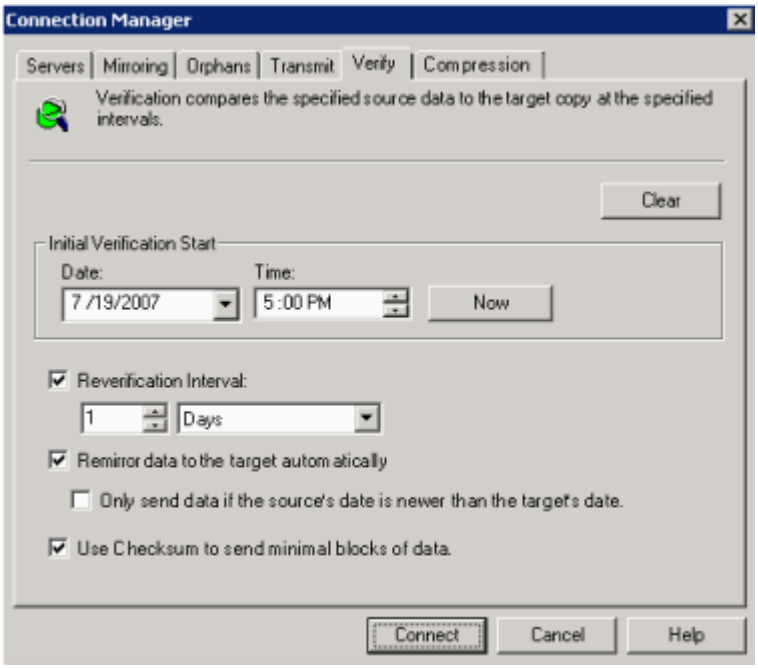
NOTE:	Database applications may update files without changing the date, time, or file size. Therefore, if you are using database applications, you should use the block checksum comparison to ensure proper verification and remirroring.

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

Scheduled verification

Verification can be scheduled to occur automatically at periodic intervals.

1. Right-click the connection on the right pane of the Management 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 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 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.

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

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.

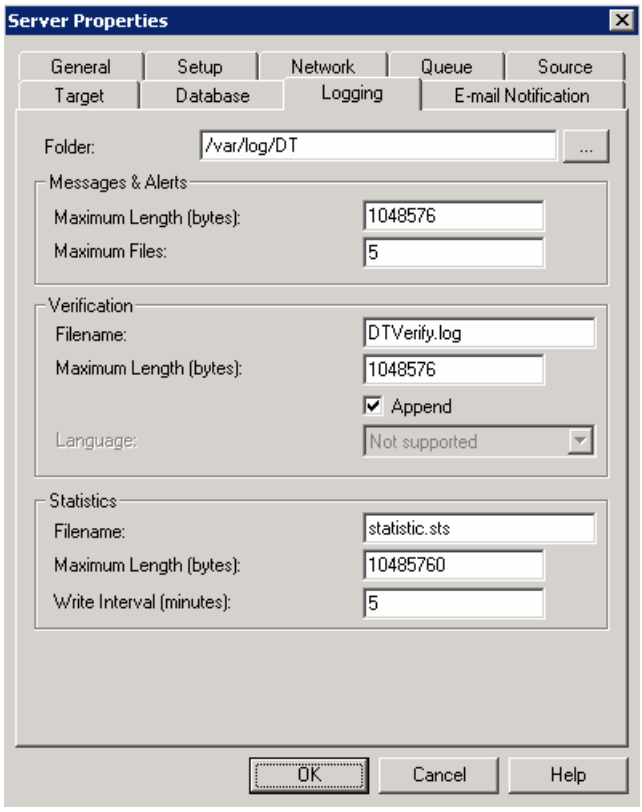
NOTE:	Database applications may update files without changing the date, time, or file size. Therefore, if you are using database applications, you should use the block checksum comparison to ensure proper verification and remirroring.

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

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 Management 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 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 (1MB). 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.

NOTE: 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.

Data Transmission 12

Double-Take 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 tranmission to suit your environment.

- **Manual transmission control**—You can maintain the source/target connection, but still control the transmission of data across the network by using the manual transmission controls. If tranmission is paused, the data is queued on the source until you manually restart the transmission.
- **Manual target control**—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.
- **Scheduled 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.
- **Bandwidth limitations**—You can specify bandwidth limitations to restrict the amount of network bandwidth used for Double-Take data transmissions. Data is queued on the source until bandwidth is available. Bandwidth limitations can be full-time or scheduled.

Controlling transmission manually

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.

Controlling a target manually

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 source cannot queue data indefinitely. If the source queue is filled, Double-Take will automatically disconnect the connections and attempt to reconnect them. See ***Configuring automatic reconnections*** on page 9-12 for more information.

To pause a target, right-click a target server on the left pane of the Management Console and select **Pause Target**. All active connections to that target will complete the operatings already in progress. You will see Pause Pending in the Management Console while these operations are completed. The status will update to Paused after the operations are completed. Any new operations will be queued on the source until the target is resumed. When you are ready to resume the target, right-click the target and select **Resume Target**.

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

Scheduling transmission criteria and bandwidth limiting

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

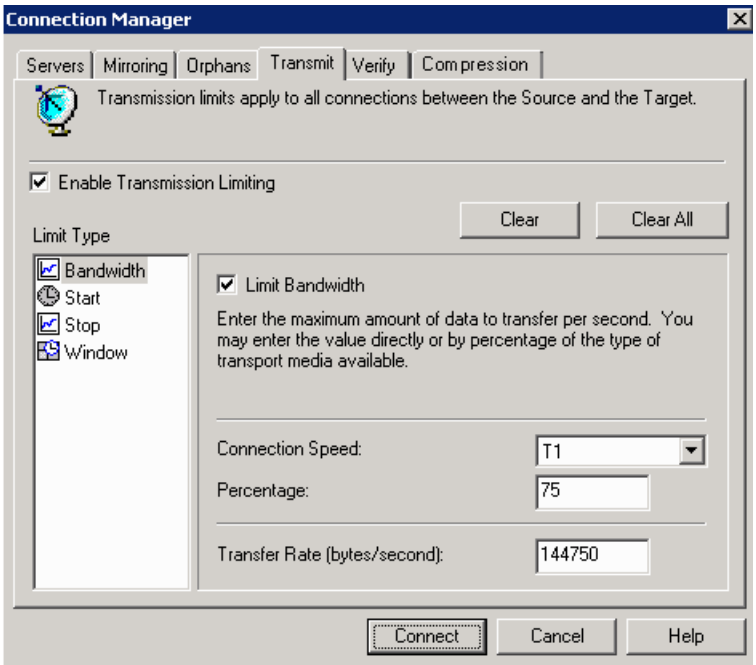
NOTE: Double-Take checks the schedule once every second, and if a user-defined criteria is met, transmission will start or stop, depending on the option specified.

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 Management 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.

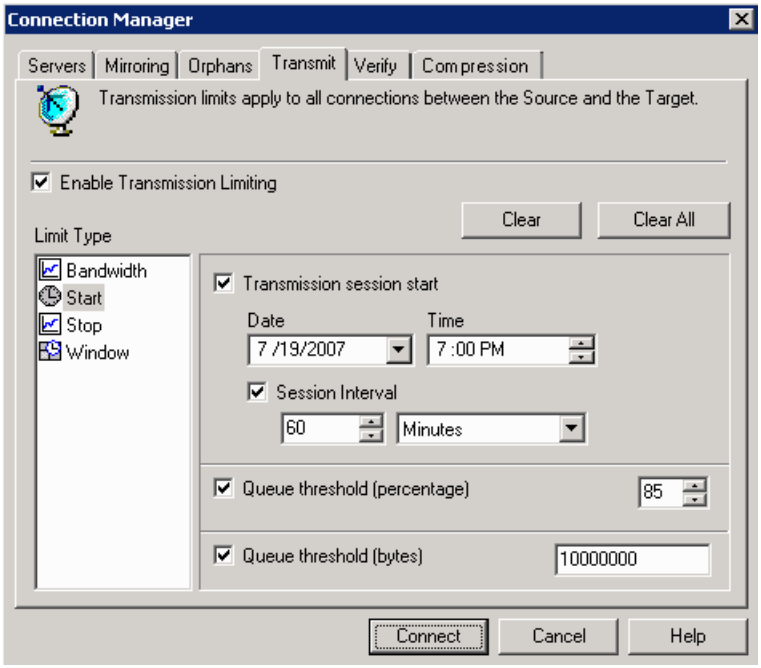
   Also at the top of the **Transmit** tab dialog box, the **Clear All** button, when selected, will remove all transmission limitations that have been set under any of the limit types. The **Clear** button will clear the settings only for the **Limit Type** selected.
3. Select the **Bandwidth** option in the **Limit Type** box. Mark the **Limit Bandwidth** check box to enable the bandwidth limiting features. Define the bandwidth available for Double-Take transmission by using either of the following options.



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

NOTE: The only value that is persistently stored is the number of bytes per second. When the page is refreshed, the percentage and available bandwidth may not be the same value that you entered. Double-Take changes these values to the maximum values for the smallest possible link. For example, if you enter 10% of T1, this is equivalent to 19300 bytes/second. If you select another transmit option or another tab of the Connection Manager and then return to the bandwidth option, Double-Take will change the value to 60% of 256 Kbps which is also equivalent to 19300 bytes/second.

4. When you schedule transmission start criteria, transmission will start when the criteria is met and will continue until the queue is empty or a transmission stop criteria is met. Select the **Start** option in the **Limit Type** box.

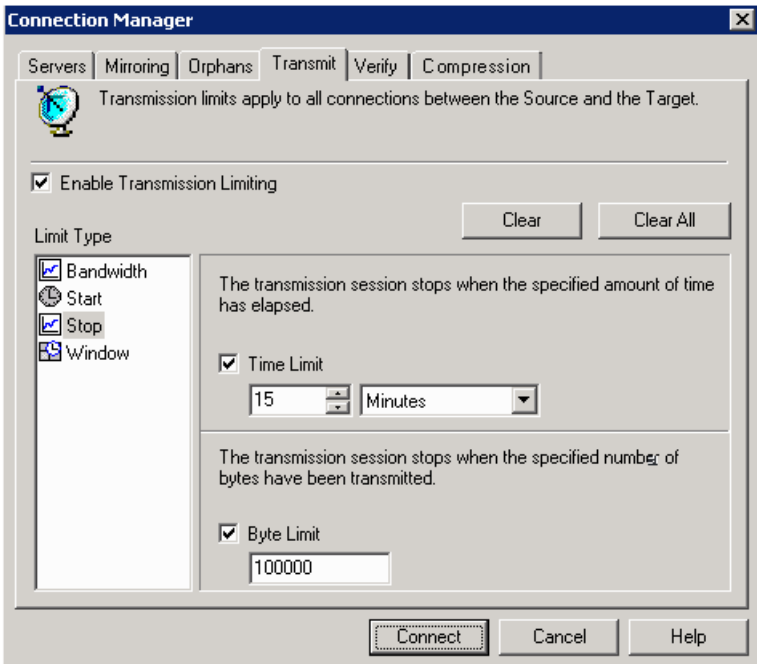


Define the start options for Double-Take transmission by using any combination of the following options.

- **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 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 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 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 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 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 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 stop transmission criteria is met. Or you might set a queue threshold of 500 MB. Double-Take will wait until there is 500 MB of data in the queue and then begin transmitting the data. Like other start criteria, Double-Take continues transmitting until the queue is empty or a Double-Take stop criteria is met. Specify a percentage of the disk queue and system memory that must be in use to initiate the transmission process, and/or specify the number of bytes that must be in the source queue and system memory to initiate the transmission process.

NOTE:	A **Transmission Session Start** setting will override any other start criteria. For example, if you set the **Transmission Session Start** and the **Queue Threshold**, transmission will not start until you reach the indicated start time.

5. Schedule any desired stop criteria to stop transmission after a transmission start criteria has initiated the transmission. If you do not establish a stop criteria, transmission will end when the queue is empty. Select the **Stop** option in the **Limit Type** box.



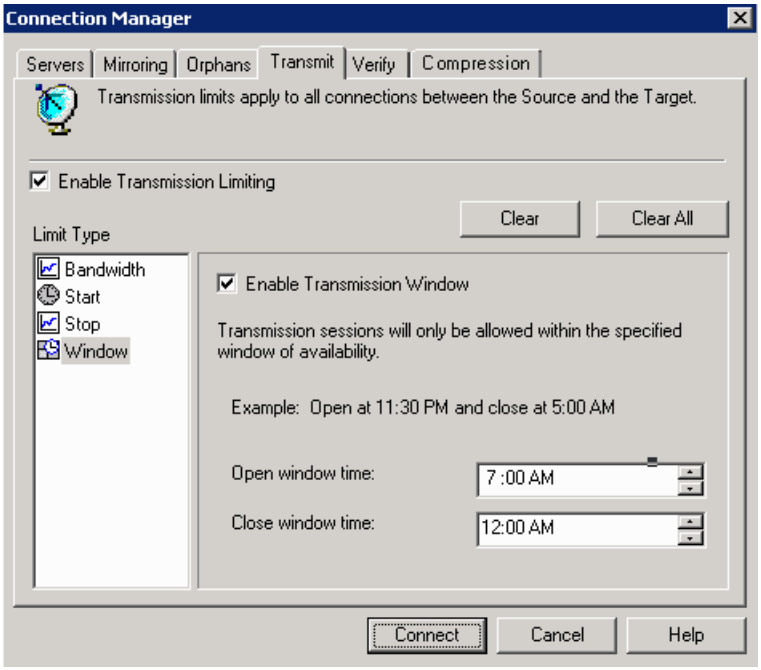
Define the stop options to stop Double-Take transmissions by using either or both of the following options.

- **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 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 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 transmission.

NOTE:	The transmission start and stop criteria should be used in conjunction with each other. For example, if you set the Queue Threshold equal to 10 MB and the Byte Limit equal to 10 MB, a network connection will be established when there is 10 MB of data in the queue. The data will be transmitted and when the 10 MB Byte Limit is reached, the network connection closes. This is useful in configurations where metered charges are based on connection time.

6. Schedule a transmission window to establish a period of availability for all Double-Take transmissions. You can specify a begin and end time for all Double-Take transmissions. When a transmission window is in effect, all other start and stop criteria are bound by this window. This means that Double-Take will never transmit data outside of an established window, regardless of other transmission settings. For example, if you set a window of availability from 9 p.m. to 4 a.m. and a start option to initiate transmission at 5 a.m., the window option will override the start option and no data will be sent at 5 a.m. Select the **Window** option in the **Limit Type** box.

NOTE: Setting a transmission window by itself is not sufficient to start a transmission. You still need to set a start criteria within the window.



Define a window to control Double-Take transmissions by enabling the feature and then specifying both window options.

- **Enable Transmission Window**—This option specifies whether a transmission window is in use.
- **Open window time**—Specifies the time, formatted for AM or PM, when the transmission window will open, allowing transmission to begin.
- **Close window time**—Specifies the time, formatted for AM or PM, when the transmission window will close, stopping all transmission.

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

Failover and Failback 13

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 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 assumes that the machine has failed. Double-Take 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 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.

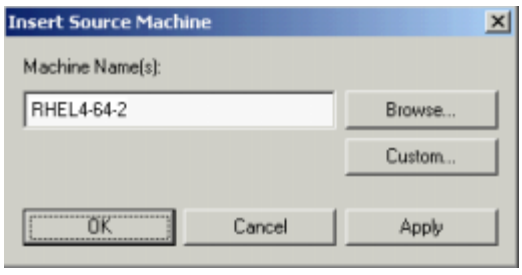
Failover can be configured to stand in for one or more IP addresses associated with different NICs on the source. Each IP address can be added to a specific target NIC making NIC configuration very flexible. For example, a single NIC on the source may have one or more IP addresses assigned to it. If 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.

Configuring failover

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

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

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



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

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

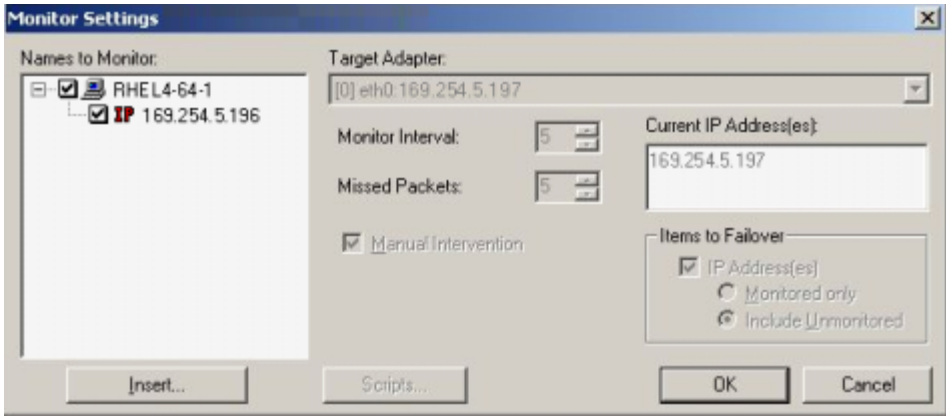
The Insert Source Machine dialog closes and the Monitor Settings dialog remains open with your source listed in the **Names to Monitor** tree.

6. In the **Names to Monitor** tree, locate and select the IP addresses on the source that you want to monitor.
7. Highlight an IP address that you have selected for monitoring and select a **Target Adapter** that will assume that IP address during failover. Repeat this process for each IP address that is being monitored.

NOTE: **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.

NOTE: 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.
      - **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.

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

If your network environment is a WAN configuration, do not failover your IP addresses unless you have a VPN infrastrucutre so that the source and target can be on the same subnet, in which case IP address failover will work the same as a LAN configuration. If you do not have a VPN, you can automatically reconfigure the routers via a failover script (by moving the source's subnet from the source's physical network to the target's physical network). There are a number of issues to consider when designing a solution that requires router configuration to achieve IP address failover. Since the route to the source's subnet will be changed at failover, the source server must be the only system on that subnet, which in turn requires all server communications to pass through a router. Additionally, it may take several minutes or even hours for routing tables on other routers throughout the network to converge.

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.

    - If you want to delay the failover or failback processes until the associated script has completed, mark the appropriate check box.
    - If you want the same scripts to be used as the default for future monitor sessions, mark the appropriate check box.

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

Editing failover 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*** on page 13-1.

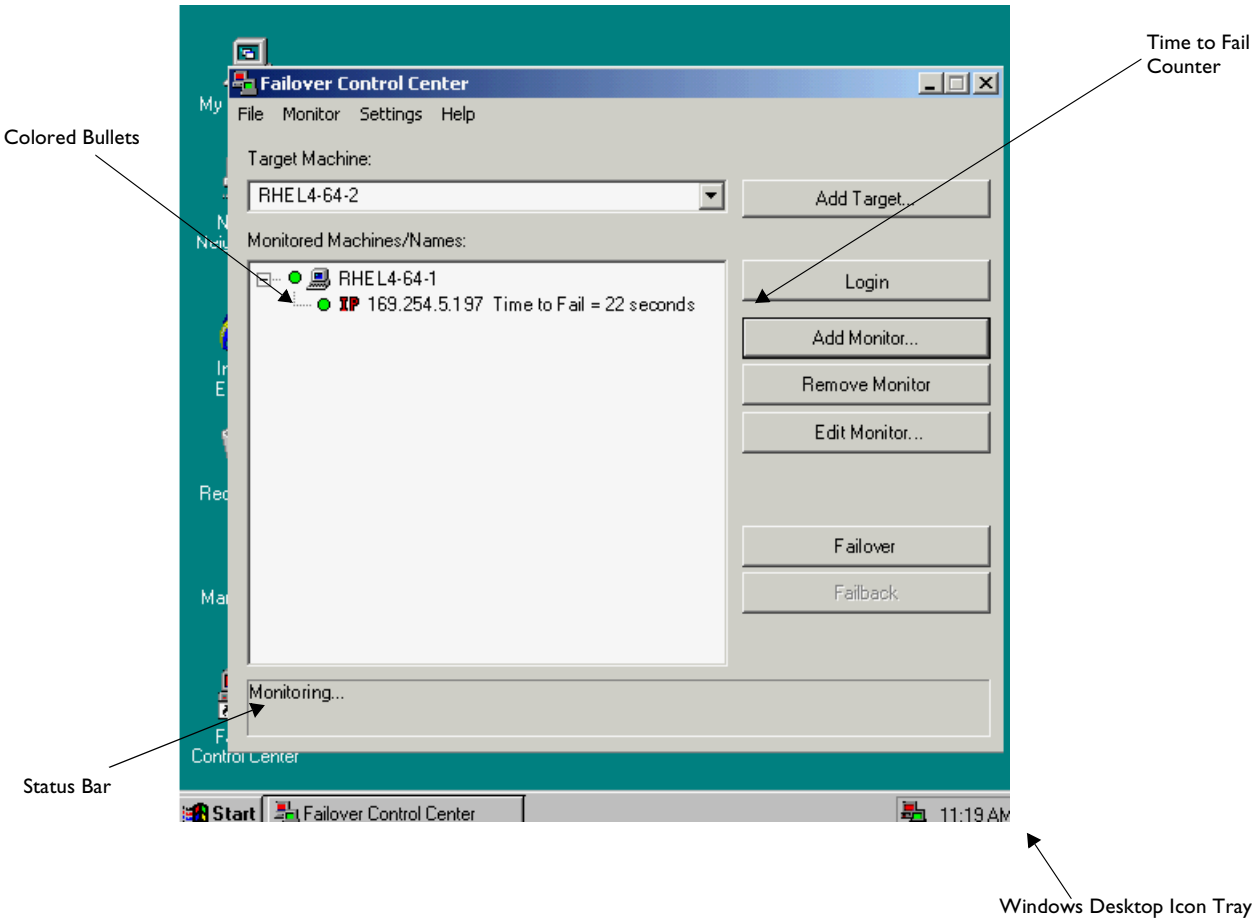
Removing failover configuration

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.

Monitoring failover

Since it can be essential to quickly know the status of failover, Double-Take offers various methods for monitoring the state of failover. When the Failover Control Center is running, you will see four visual indicators:

- The Failover Control Center Time to Fail counter
- The Failover Control Center status bar located at the bottom of the window
- The Failover Control Center colored bullets to the left of each IP address and source machine
- The Windows desktop icon tray containing a failover icon



NOTE: You can minimize the Failover Control Center and, although it will not appear in your Windows taskbar, it will still be active and the failover icon will still appear in the desktop icon tray.

The Failover Control Center does not have to be running for failover to occur.

The following table identifies how the visual indicators change as the status of failover changes.

	Time to Fail Countdown	Status Bar	Colored Bullets	Desktop Icon Tray
Source is Online	The Time to Fail counter is counting down and resetting each time a response is received from the source machine.	The status bar indicates that the target machine is monitoring the source machine.	The bullets are green. [a]	The Windows desktop icon tray contains a failover icon with red and green computers.
Source Fails and Failover is Initiated	The Time to Fail countdown value is 0.	The status bar displays the source machine and IP address currently being assumed by the target.	The bullets are red.	The Windows desktop icon tray contains a failover icon with red and green computers.
Failover is Complete	The Time to Fail counter is replaced with failed message.	The status bar indicates that monitoring has continued.	The bullets are red.	The Windows desktop icon tray contains a failover icon with a red computer.

a. 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.

Testing failover

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**.

Failover manual intervention

If **Manual Intervention** is enabled, the Failover Control Center will prompt you when a failure occurs.

NOTE: If the Failover Control Center is not running at the time the failure occurs, the manual intervention dialog box will appear the next time the Failover Control Center is started.

When a failure occurs, an alert is forwarded to the 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.

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.

Initiating failback

When failover occurs, a source machine has failed. The steps below must be completed in order to complete failback without IP address or name conflicts.

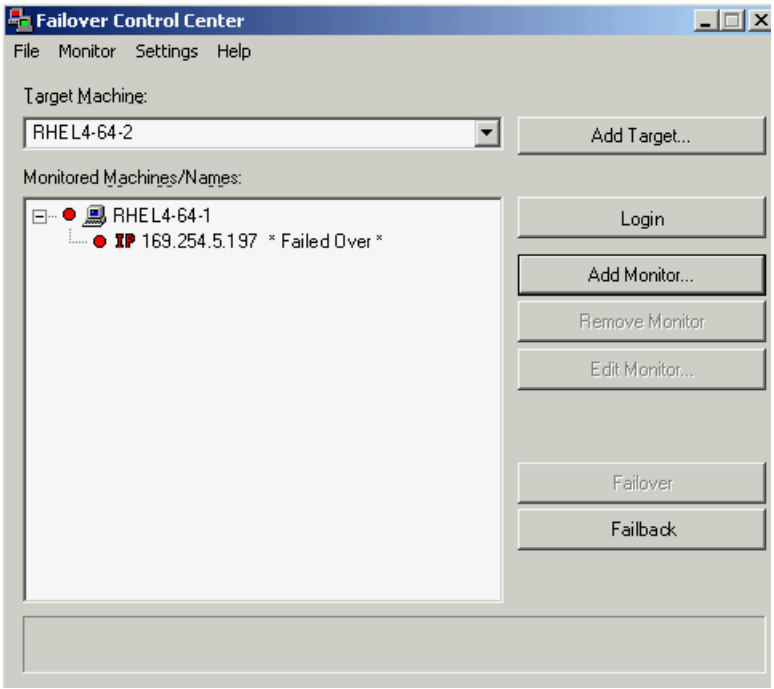
1. Verify that your source machine is not connected to the network. If it is, disconnect it.
2. Resolve the source machine problem that caused the failure.

WARNING:	Do not connect the source machine to the network at this time.

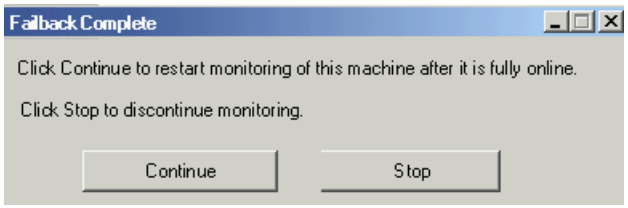
3. In the Failover Control Center, select the source machine that is currently standing in for the failed source.

NOTE:	If your target machine has been replaced during failover because of the configuration you selected (perform failover by replacing target identity with source), do not select your target from the Target Machine list. Since the target has been replaced by the source, it no longer exists and the Failover Control Center will retry for several minutes (possibly up to five minutes), until it finally fails in locating the target. After a replace failover, select your source machine in the Target Machine list, which will be monitoring itself in a failed over state, and perform failback. After failback, the target will be available again.

4. Select the failed source and click Failback. If you specified a pre-failback script in your failover configuration, that script will be executed at this time.



5. You will be prompted to determine if you want to continue monitoring the source. Bring the source online and 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.

NOTE:	At this time, the source machine must be online and Double-Take 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.

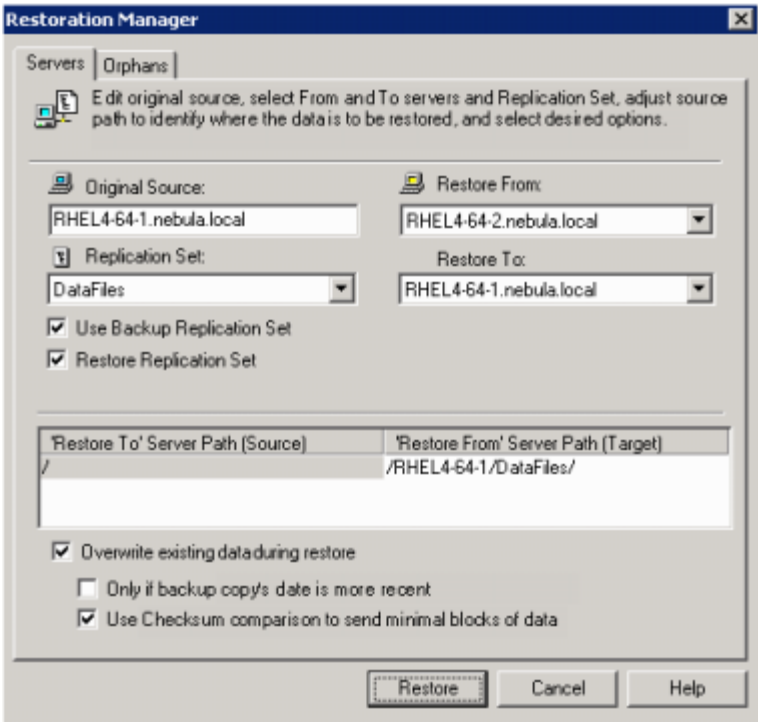
Failback is now complete.

Restoration

14

The Management 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 handles the rest, including selecting the files in the replication set and restoring them to the correct location.

1. From the Management 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.

NOTE:	Restoration applies to the entire replication set and should not be used for individual file restoration. Individual files can be restored manually by copying the files from the target to the source machine using native operating system tools.

5. Select the **Restore To** machine. 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.

NOTE:	Restoring across a NAT router requires the ports to be the same as the original connection. If the ports have been modified (manually or reinstalled), you must set the port numbers to the same values as the last valid source/target connection.

7. 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.
8. 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.
9. 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**—This option restores only those files that are newer on the target than on the source. The entire file is overwritten with this option.

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

   - **Use 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.
10. 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. See ***Configuring and removing orphan files*** on page 10-5 for more details.
11. Click **Restore** to begin the restoration.

Monitoring Tools

15

In addition to the monitoring capabilities within the Double-Take clients, you can monitor Double-Take processing using other tools.

- Log files are generated for various Double-Take components recording alerts (notifications, warnings, and errors). These log files can be viewed using a standard text viewer. See ***Log files*** on page 15-2.
- The Linux system log also receives alerts (notifications, warnings, and errors), but only from the Double-Take daemon. These alerts are logged to the application log. See ***Linux System Log*** on page 15-9.
- DTStat is a utility that reports statistical information. See ***DTStat*** on page 15-19.
- SNMP reports alerts and statistics. See ***SNMP*** on page 15-26.

Log files

Various Double-Take components (Double-Take daemon, Management 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. They can be viewed, from the location where Double-Take is installed, with a standard text viewer.

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.

Component	Log File Base Name
Double-Take	dtlog
Management Console	mc
Failover Control Center	fcc
Command Line Client	dtcl

- **Series Number**—The series number ranges from `1` to `999`. For example, Double-Take 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 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`.

Component	Sample Log File Names
Double-Take	dtlog1.dtl, dtlog2.dtl
Management Console	mc1.dtl, mc2.dtl
Failover Control Center	fcc1.dtl, fcc2.dtl
Command Line Client	dtcl1.dtl, dtcl2.dtl

1. To modify the maximum file size and the number of Double-Take log files that are maintained, access the Server Properties dialog box by right-clicking a machine name in the left pane of the Management 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 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`.

NOTE:
- 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.

Viewing the log files

The log files can be viewed, from the location where Double-Take is installed, with a standard text viewer.

Sample Double-Take Daemon Log File

```
01/15/2007 14:14:18.3900   95   98   2 2      69  Kernel Started
01/15/2007 14:14:18.4200   95   98   3 2   10004  Valid Activation Key Detected :
01/15/2007 14:14:18.5350   98 170   4 2   52501  Target module loaded successfully
01/15/2007 14:14:18.6760   98 172   5 2   10004  Valid Activation Key Detected :
01/15/2007 14:14:18.9870 130 131   6 2   51501  Source module loaded successfully
01/15/2007 14:24:15.2070 130 132   7 2      72  Connection Request from
ip://206.31.4.305
01/15/2007 14:24:16.3090 131 133   8 2 600002 Unified login provides ADMIN access
01/15/2007 14:24:40.9680 132 134   9 2      99   RepSet Modified: UserData
01/15/2007 14:25:22.4070 134 131 10 2    71   Originator Attempting ip://206.31.4.305
01/15/2007 14:25:22.5030 134 131 11 2       0   Transmission Create to
ip://206.31.4.305.
01/15/2007 14:25:22.6060 135 133 12 2 500000   UserData is connected to
ip://206.31.4.305
01/15/2007 14:25:23.5030 136   98 13 2    87 Start Replication on connection 1
```

Sample Management Console Log File

```
00/00/0000  00:00:00.0000 Application starting
09/11/2007  12:45:53.8980 704 1032 1 2 0  Could not find XML file: C:\Program
Files\DoubleTake\ Administrator.xml, default groups will be added.
09/11/2007  12:45:53.9580 704 1032 2 2 0  Adding default group: Double-Take Servers
09/11/2007  12:45:53.9580 704 1032 3 2 0  Adding default group: Double-Take Servers\
Auto-Discovered Servers
09/11/2007  12:46:08.3390 704 1032 4 210004  Evaluation license expires in 95 day(s).
```

The following table describes the information found in each column of the log file.

Column #	Description
Column 1	The date the message was generated.
Column 2	The time the message was generated.
Column 3	The process ID.
Column 4	The thread ID.
Column 5	The sequence number is an incremental counter that assigns a unique number to each message.
Column 6	The type or level of message displayed: **1**—Warning or error message **2**—Informational message
Column 7	The message ID.
Column 8	The message text.

Log messages

The following table describes some of the standard Double-Take 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.

NOTE: In the following table, *con_id* refers to the unique connection ID assigned to each connection between a source replication set and a target.

ID	Message	Description
0	N/A	There are several log messages with this ID number. See the description in the Message column.

7	• Synchronous ioctl returned STATUS_ PENDING
	• Failed to reset Replication Flags. Replication may not be performed correctly.
		• Communication with the Double-Take driver is not being performed correctly. A reboot is required to guarantee replication and data integrity.
		• An error occurred between the Double-Take driver and recent changes to the replication set. The possible resolutions are:
		  • Undo the changes to the replication set
		  • Stop and restart Double-Take
		  • Reboot the server

69	Double-Take kernel started on *server_name*	The Double-Take daemon was started on the Double-Take server specified.
70	Double-Take kernel stopped	The Double-Take daemon was stopped on a Double-Take 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 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.

ID	Message	Description
78	Auto-disconnect threshold has been reached.	The Double-Take 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 Management 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 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.
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.
85	*repset* has been auto-disconnected	Double-Take 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 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 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 data on the target is identical to the data on the source has ended.

ID	Message	Description
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 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 *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 daemon 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 is disabled.
10004	Valid activation key detected	A valid activation code was identified when the Double-Take daemon was started.
51001	Source module failed to load	The Double-Take 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 source module was loaded successfully.
51502	Source module already loaded	The Double-Take source module was already loaded.
51503	Source module stopped	The Double-Take source module stopped.
52000	• The target has been paused due to manual intervention. • The target has been resumed due to manual intervention.	The target has been paused or resumed through user intervention.

ID	Message	Description
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. The most common code values that appear in this error message are: • **(5) Permission denied**: The account running the Double-Take daemon does not have permission to update the file specified. • **(32) Sharing violation:** Another application is using a particular file that Double-Take is trying to update. Double-Take 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 target module was loaded successfully.
52502	Target module already loaded	The Double-Take target module was already loaded.
52503	Target module stopped	The Double-Take 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 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 daemon was started.
54001	Failover module failed to load	The Double-Take failover module failed to load. Look at previous log messages to determine the reason.
54503	Failover module stopped	The Double-Take 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 Management Console.
99999	Application is terminating normally	The Double-Take daemon is shutting down normally.
503010	Asyncloctl for status thread 178 terminated, terminating the status thread	A Double-Take process monitors the state of the Double-Take driver. When the daemon is shut down, the driver is shut down, and this process is terminated. (No user action required.)

ID	Message	Description
600002	• Unified login provides ADMIN access • User *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 occurence and will continue until the source machine responds or until failover occurs.
800000	• Active Directory GetHostSpns function call failed • Active Directory RemoveSpns function call failed • Active Directory AddSpns function call failed	• Double-Take failed to get the host SPN (Service Principal Name) from Active Directory. • Double-Take failed to remove an SPN from Active Directory. • Double-Take failed to add a host SPN to Active Directory.

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.

ID	Category	Severity	Event Message	Required Response
1	Activation Key	Error	This evaluation period has expired. Mirroring and replication have been stopped. To obtain a license, please contact your vendor.	Contact your vendor to purchase either a single or site license.
2	Activation Key	Info.	The evaluation period expires in %1 day(s).	Contact your vendor before the evaluation period expires to purchase either a single or site license.
3	Activation Key	Info.	The evaluation period has been activated and expires in %1 day(s).	Contact your vendor before the evaluation period expires to purchase either a single or site license.
4	Activation Key	Warning	Duplicate activation codes detected on machine %1 from machine %2.	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	Activation Key	Error	This product edition can only be run on Windows 2000 Server or Advanced Server running the Server Appliance Kit.	Verify your activation code has been entered correctly and contact technical support.
3000	Service	Info.	Logger service was successfully started.	No action required.
3001	Service	Info.	Logger service was successfully stopped.	No action required.
4000	Service	Info.	Kernel was successfully started.	No action required.
4001	Service	Info.	Target service was successfully started.	No action required.
4002	Service	Info.	Source service was successfully started.	No action required.
4003	Service	Info.	Source service was successfully stopped.	No action required.
4004	Service	Info.	Target service was successfully stopped.	No action required.
4005	Service	Info.	Kernel was successfully stopped.	No action required.
4006	Service	Error	Service has aborted due to the following unrecoverable error: %1	Restart the Double-Take daemon.

ID	Category	Severity	Event Message	Required Response
4007	Service	Warning	Auto-disconnecting from %1 (%2) for Replication Set %3, ID: %4 due to %5	The connection is auto-disconnecting because the disk-based queue on the source has been filled, the daemon 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 cannot read from or write to the transaction log file).
4008	Service	Info.	Auto-disconnect has succeeded for %1 (%2) for Replication Set %3, ID: %4	No action required.
4009	Service	Info.	Auto-reconnecting Replication Set %1 to %2 (%3)	No action required.
4010	Service	Info.	Auto-reconnect has succeeded connecting Replication Set %1 to %2 (%3)	No action required.
4011	Service	Error	Auto-reconnect has failed connecting Replication Set %1 to %2 (%3)	Manually reestablish the replication set to target connection.
4014	Service	Info.	Service has started network transmission.	No action required.
4015	Service	Info.	Service has stopped network transmission.	No action required.
4016	Service	Info.	Service has established a connection to %1 (%2) for Replication Set %3, ID: %4	No action required.
4017	Service	Info.	Service has disconnected from %1 (%2) for Replication Set %3, ID: %4	No action required.
4018	Service	Warning	%1, however, mirroring and replication have been disabled as a restore is required due to a previous failover.	Perform a restoration.
4019	Service	Info.	Service has started a mirror to %1 (%2) for Replication Set %3, ID: %4	No action required.
4020	Service	Info.	Service has paused a mirror to %1 (%2) for Replication Set %3, ID: %4	No action required.
4021	Service	Info.	Service has resumed a mirror to %1 (%2) for Replication Set %3, ID: %4	No action required.
4022	Service	Info.	Service has stopped a mirror to %1 for Replication Set %2, ID: %3, %4	No action required.

ID	Category	Severity	Event Message	Required Response
4023	Service	Info.	Service has completed a mirror to %1 %2 for Replication Set %3, ID: %4, %5	No action required.
4024	Service	Info.	Service has started Replication to %1 (%2) for Replication Set %3, ID: %4	No action required.
4025	Service	Info.	Service has stopped Replication to %1 (%2) for Replication Set %3, ID: %4	No action required.
4026	Service	Info.	The target has been paused due to user intervention.	No action required.
4027	Service	Info.	The target has been resumed due to user intervention.	No action required.
4028	Service	Warning	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.	Verify that the Active Directory server is running and that the Double-Take daemon has permission to update Active Directory.
4029	Service	Warning	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.	Verify that the Active Directory server is running and that the Double-Take daemon has permission to update Active Directory.
4030	Service	Error	RSResource.dll has an unknown error. The product functionality has been disabled.	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	Service	Error	RSResource.dll could not be opened. The product functionality has been disabled.	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	Service	Error	The RSResource.dll component version does not match the component version expected by the product. The product functionality has been disabled.	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	Service	Error	RSResource.dll build version is invalid. The product functionality has been disabled.	Reinstall the software, using the installation Repair option, to install a new copy of the RSResource.dll. Contact technical support if this error persists.

ID	Category	Severity	Event Message	Required Response
4034	Service	Error	Error verifying the service name. The product functionality has been disabled.	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	Service	Error	Error verifying the product name. The product functionality has been disabled.	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	Service	Error	Error verifying the vendor name. The product functionality has been disabled.	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	Service	Error	Error verifying the vendor URL name. The product functionality has been disabled.	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	Service	Error	Error verifying the product code. The product functionality has been disabled.	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	Service	Error	Error while reading RSResource.dll. The product functionality has been disabled.	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	Service	Error	The product code is illegal for this computer hardware. The product functionality has been disabled.	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	Service	Error	The product code is illegal for this operating system version. The product functionality has been disabled.	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	Service	Error	The product code requires installing the Windows Server Appliance Kit. The product functionality has been disabled.	Reinstall the software, using the installation Repair option, to install a new copy of the RSResource.dll. Contact technical support if this error persists.

ID	Category	Severity	Event Message	Required Response
4043	Service	Error	This product can only be run on a limited number of processors and this server exceeds the limit. The product functionality has been disabled.	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	Service	Error	An error was encountered and replication has been stopped. It is necessary to stop and restart the service to correct this error.	Contact technical support if this error persists.
4045	Service	Error	%1 value must be between 1025 and 65535. Using default of %2.	Verify that the Double-Take 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	Service	Error	This service failed to start because of a possible port conflict. Win32 error: %1	Verify that the Double-Take ports are not conflicting with ports used by other applications.
4047	Service	Error	Could not load ZLIB DLL %1. Some levels of compression will not be available.	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	Info.	Service has started a delete orphans task to %1 (%2) for Replication Set %3, ID: %4	No action required.
4049	Service	Info.	Service has paused a delete orphans task to %1 (%2) for Replication Set %3, ID: %4	No action required.
4050	Service	Info.	Service has resumed a delete orphans task to %1 (%2) for Replication Set %3, ID: %4	No action required.
4051	Service	Info.	Service has stopped a delete orphans task to %1 (%2) for Replication Set %3, ID: %4	No action required.
4052	Service	Info.	Service has completed a delete orphans task to %1 (%2) for Replication Set %3, ID: %4	No action required.
4053	Service	Info.	Service has started a restore task to %1 (%2) for Replication Set %3, ID: %4	No action required.
4054	Service	Info.	Service has paused a restore task to %1 (%2) for Replication Set %3, ID: %4	No action required.
4055	Service	Info.	Service has resumed a restore task to %1 (%2) for Replication Set %3, ID: %4	No action required.

ID	Category	Severity	Event Message	Required Response
4056	Service	Info.	Service has stopped a restore task to %1 (%2) for Replication Set %3, ID: %4	No action required.
4057	Service	Info.	Service has completed a restore task to %1 (%2) for Replication Set %3, ID: %4	No action required.
4058	Service	Info.	Service has started a verification task to %1 (%2) for Replication Set %3, ID: %4	No action required.
4059	Service	Info.	Service has paused a verification task to %1 (%2) for Replication Set %3, ID: %4	No action required.
4060	Service	Info.	Service has resumed a verification task to %1 (%2) for Replication Set %3, ID: %4	No action required.
4061	Service	Info.	Service has stopped a verification task to %1 (%2) for Replication Set %3, ID: %4	No action required.
4062	Service	Info.	Service has completed a verification task to %1 (%2) for Replication Set %3, ID: %4	No action required.
4110	Service	Warning	Target cannot write %1 due to target disk being full. Operation will be retried (%2 times or forever)	The disk on the target is full. The operation will be retried according to the TGExecutionRetryLimit setting.
4111	Service	Warning	Target can not write %1 due to a sharing violation. Operation will be retried (%2 times or forever)	A sharing violation error is prohibiting Double-Take from writing on the target. The operation will be retried according to the TGExecutionRetryLimit setting.
4112	Service	Warning	Target can not write %1 due to access denied. Operation will be retried (%2 times or forever)	An access denied error is prohibiting Double-Take from writing on the target. The operation will be retried according to the TGExecutionRetryLimit setting.
4113	Service	Warning	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.	An unknown error is prohibiting Double-Take from writing on the target. The operation will be retried according to the TGExecutionRetryLimit setting.
4120	Service	Info.	Target write to %1 was completed successfully after %2 retries.	No action required.

ID	Category	Severity	Event Message	Required Response
4150	Service	Error	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.	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	Service	Info.	In band task %1 submitted from %2 by %3 at %4	No action required.
4201	Service	Warning	In band task %1 discarded (submitted from %2 by %3 at %4)	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	Service	Info.	Running %1 in band script: %2 (task %3 submitted from %4 by %5 at %6)	No action required.
4203	Service	Info.	Completed run of in band script: %1 (exit code %2)	No action required.
4204	Service	Error	Error running in band script: %1	Review the task and its associated script(s) for syntax errors.
4205	Service	Warning	Timeout (%1 seconds) running in band script: %2	The timeout specified for the script to complete has expired. Double-Take will continue with normal processing. You may need to manually terminate the script if it will never complete.
4206	Service	Warning	Run timeout disabled for in band script: %1	The timeout period was set to zero (0). Double-Take will not wait for the script to complete before continuing. No action is required.
4207	Service	Warning	In band scripts disabled by server - no attempt will be made to run %1	Enable task command processing.
4300	Service	Error	A connection request was received on the target before the persistent target paths could be loaded.	You may need to disconnect and reconnect your replication set.
4301	Service	Error	Unable to block target paths, the driver is unavailable.	If you need to block your target paths, contact technical support.
4302	Service	Info.	Target Path %1 has been successfully blocked	No action required.

ID	Category	Severity	Event Message	Required Response
4303	Service	Warning	Blocking of target path: %1 failed. Error Code: %2	If you need to block your target paths, contact technical support.
4304	Service	Info.	Target Path %1 has been successfully unblocked	No action required.
4305	Service	Warning	Unblocking of target path: %1 failed. Error Code: %2	If you need to unblock your target paths, contact technical support.
4306	Service	Warning	Target paths for source %1 (%2) Connection id: %3 are already blocked	No action required.
4307	Service	Warning	Target paths for source %1 (%2) Connection id: %3 are already unblocked	No action required.
4308	Service	Error	Error loading target paths for blocking, registry key %1 has been corrupted.	If you need to block your target paths, contact technical support.
5000	Service	Info.	Server Monitor service was successfully started.	No action required.
5001	Service	Info.	Server Monitor service was successfully stopped.	No action required.
5002	Service	Info.	Placeholders were modified to %1.	No action required.
5100	Failover	Info.	Failover completed for %1.	No action required.
5101	Failover	Info.	IP address %1 with subnet mask %2 was added to target machine's %3 adapter.	No action required.
5102	Failover	Warning	%1 has reached a failover condition. A response from the user is required before failover can take place.	User intervention has been configured. Open the Failover Control Center and accept or decline the failover prompt.
5103	Failover	Info.	Started adding drive shares from %1 to %2.	No action required.
5104	Failover	Info.	%1 drive shares were taken over by %2.	No action required.
5105	Failover	Info.	Attempting to run the %1 script.	No action required.
5106	Failover	Info.	The %1 script ran successfully.	No action required.
5107	Failover	Error	Error occurred in running %1 script.	Verify that the script identified exists with the proper permissions.

ID	Category	Severity	Event Message	Required Response
5108	Failover	Error	The source machine %1 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.
5109	Failover	Error	The public NIC on source machine %1 is not responding to a ping.	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	Info..	Failback completed for %1.	No action required.
5201	Failback	Info..	IP address %1 was removed from target machine's %2 adapter.	No action required.
5202	Failback	Error	Unable to Failback properly because IP address %1 was missing a corresponding SubNet Mask.	Contact technical support.
5300	Monitoring	Info.	The following IP address was added to target's monitoring list: %1	No action required.
5301	Monitoring	Info.	The following IP address was removed from target's monitoring list: %1	No action required.
5302	Monitoring	Info.	Drive share information for %1 has been updated on the target machine.	No action required.
5500	Service	Warning	Could not connect to e-mail server. Check to make sure the SMTP server %1 is available (error code: %2).	Double-Take 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	Service	Warning	E-mail notification could not be enabled (error code: %1).	This alert occurs if there is an unexpected error enabling e-mail notification during daemon 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) daemon is enabled. If neither of these apply, contact technical support.

ID	Category	Severity	Event Message	Required Response
5502	Service	Warning	E-mail notification could not be initialized. Check to make sure Internet Explorer 5.0 or later is installed.	E-mail notification requires Internet Explorer 5.0 or later. Verify that you have this version or later installed on the Double-Take server.
5503	Service	Warning	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).	If you are using Double-Take 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 server.
5504	Service	Warning	Could not load LocalRS.dll (for e-mail notification).	This alert occurs if there is an error loading the resource DLL for the daemon. 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	Service	Warning	E-mail could not be sent. Check e-mail settings (error code: %1).	Verify that the e-mail server that you have identified in your e-mail notification configuration is correct.
5506	Service	Warning	One or more required e-mail settings have not been specified (error code: %1).	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	Service	Warning	E-mail notification could not be initialized. Check to make sure WMI is installed and available (error code: %1).	If you are using Double-Take 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 server.
5508	Service	Warning	An error occurred connecting to the WMI namespace. Check to make sure the Windows Management Instrumentation service is not disabled (error code %1).	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 server and that it is enabled.
5600	Service	Warning	Part or all of the e-mail setting %1 is not in a valid format.	Verify that the include categories and exclude ID list are identified and formatted correctly.

DTStat

Statistics logging is the process of taking snapshots of Double-Take statistical data. The data can be written to a file for future use. Changes to the DTStat file configuration are detected and applied immediately without restarting the Double-Take daemon.

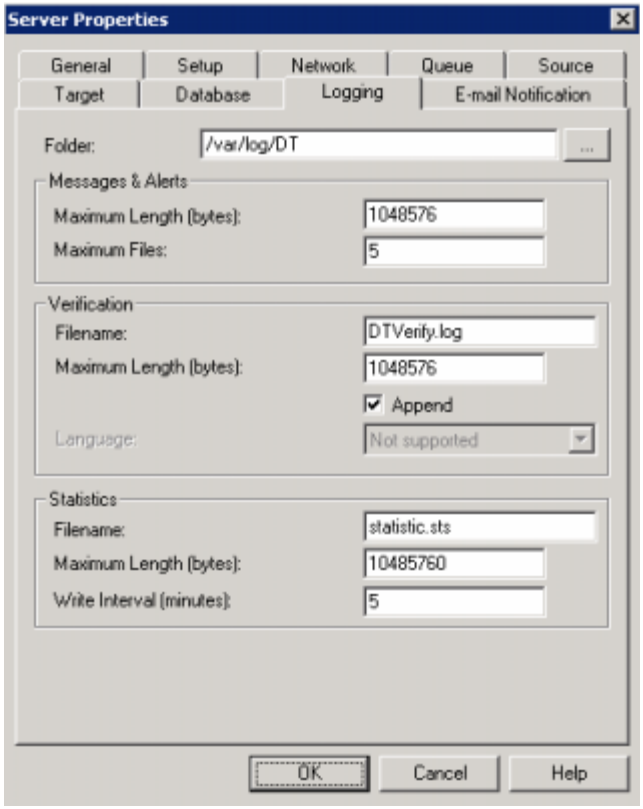
The statistics log file created is a binary file. To view the log file, you must run DTStat from the command prompt.

Sample DTStat Output

```
================================
09/11/07 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 : /var/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 statistics file

1. Right-click a machine in the left pane of the Management 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 begins overwriting the oldest data in the file.
   - **Write Interval**—The frequency in which Double-Take writes the statistical data to the statistics log file. The default is every `5` minutes.
5. Click **OK** to save the settings.

Running DTStat from the command prompt

From a command prompt, use the **DTStat** command from the location where Double-Take is installed.

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.

DTStat statistics

The following statistics appear in the default `statistic.sts` file.

NOTE: The categories you see will depend on the function of your server (source, target, or both).

If you have multiple IP addresses connected to one target server, you will see multiple Target sections for each IP address.

If you convert your statistics output to an ASCII, comma-delimited file using the `dtstat -s` option, keep in mind the following differences.

- The statistic labels will be slightly different in the ASCII file than in the following table.
- The statistics will appear in a different order in the ASCII file than in the following table.
- The statistics in the Target Category in the following table are not included in the ASCII file.
- The Kernel statistic Target Open Handles is not included in the ASCII file.
- The ASCII file contains a Managed Pagefile Alloc statistic which is no longer used.

Category	Statistic	Description
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
IQ Allocator	**Total Bytes**	The number of bytes currently allocated to the intermediate queue
Security	**Logins**	The number of successful login attempts
	FailedLogins	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
	TargetState	**0**—Target is not running **1**—Target is running
	Start Time	Date and time stamp indicating when the Double-Take daemon was loaded
	RepOps Generated	The number of replication operations generated by the file system driver. An op is a file system operation. Double-Take 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.
	RepBytes Generated	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).

Category	Statistic	Description
Kernel	**MirOps Generated**	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.
	MirBytes Generated	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.
	FailedMirror Count	The number of mirror operations that failed due to an error reading the file from the disk
	FailedRep Count	The number of replication operations that failed due to an error reading the file from the disk
	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.
	TargetOpen Handles	The number of handles currently open on the target
	DriverQueue Percent	The amount of throttling calculated as a percentage of the stop replicating limit
Target	**PeerAddress**	The IP address of the source machine
	Local Address	The IP address of the target machine.
	Ops Received	The total number of operations received by this machine as a target since the Double-Take daemon was loaded
	Mirror Ops Received	The total number of mirror operations received by this machine as a target since the Double-Take daemon was loaded. This number does not reset to zero for remirrors.
	Retries	The number of retries performed before all operations were completed
	OpsDropped	The number of operations skipped during a difference mirror. During a difference mirror, if Double-Take 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.
	Ops Remaining	The total number of operations that are left in the target queue
	Orphan Files Removed	The number of orphan files removed from the target machine
	Orphan Directories Removed	The number of orphan directories removed from the target machine
	Orphan Bytes Removed	The number of orphan bytes removed from the target machine
	Bytes In Target Queue	The number of bytes currently in the system memory queue on the target

Category	Statistic	Description
Target	**Bytes In Target Disk Queue**	The number of bytes currently in the disk queue on the target
	Tasks Succeeded	The number of task commands that have succeeded on the target
	TasksFailed	The number of task commands that have failed on the target
	Tasks Ignored	The number of task commands that have been ignored on the target
Source	**auto DisConnects**	The number of automatic disconnects since starting Double-Take. 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 daemon was started.
	autoRe Connects	The number of automatic reconnects since starting Double-Take. 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 daemon was started.
	lastFile Touched	The last filename that had a replication operation executed
Connection	**conPeer Address**	The IP address of the target machine
	connectTime	The time that this connection was established
	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. DTStat 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).
	conOpsIn CmdQueue	The number of operations waiting to be executed on the target
	conOpsInAck Queue	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.
	conOpsIn RepQueue	The number of replication operations currently waiting to be executed on the target

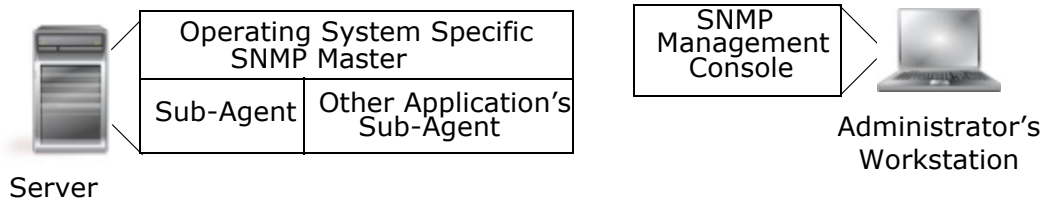
Category	Statistic	Description
Connection	**conOpsInMirQueue**	The number of mirror operations currently waiting to be executed on the target
	conBytesInRepQueue	The number of replication bytes remaining to be transmitted to the target
	conOpsTx	The number of operations transmitted to the target. This is the total number of operations that Double-Take has transmitted as a source. In other words, the cumulative number of operations transmitted by this source to all connected targets.
	conBytesInMirQueue	The number of mirror bytes remaining to be transmitted to the target
	conBytesTx	The number of bytes transmitted to the target. This is the total number of bytes that Double-Take has transmitted as a source. In other words, the cumulative number of bytes transmitted by this source to all connected targets.
	conBytesCompressedTx	The number of compressed bytes transmitted to the target.
	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.
	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.
	conResentOpCount	The number of operations resent because they were not acknowledged
	conBytesInDiskQueue	The number of bytes in the source disk queue
	conBandwidthLimit	The amount of bandwidth that may be used to transfer data
	conBytesSkipped	The number of bytes skipped during a difference mirror. During a difference mirror, if Double-Take 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.
	conMirrorBytesRemain	The number of mirror bytes remaining to be transmitted
	conMirrorPercent	The percentage of the mirror that has been completed. This field is determined if the replication set size was calculated.
	conTaskCmdsSubmitted	The number of task commands that have been submitted on the source
	conTaskCmdsQueued	The number of task commands that have been queued on the source
	conTasksSucceeded	The number of task commands that have succeeded on the source
	conTasksFailed	The number of task commands that have failed on the source
	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 provides an SNMP sub-agent that monitors Double-Take and can be managed from an SNMP Management Console.

Double-Take installs two components to work with SNMP.

- The sub-agent is a program that installs and runs on the same machine as Double-Take 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 SNMP sub-agent is included in the Double-Take installation program.
- A Double-Take 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 .mib file is NSI-DT.mib and meets SNMP standards.



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

Configuring SNMP

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 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 MIB, add the path to the Double-Take 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 SNMP sub-agent.

   #DTSubAgent >& /dev/null &

   NOTE: Instead of starting the master agent and Double-Take 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 statistics

The following tables identify the Double-Take statistics provided through SNMP.

Object Type	Statistic	Description
General	**dtUpTime**	Time in seconds since Double-Take was last started
	dtCurrentMemory Usage	Amount of memory allocated from the Double-Take memory pool
	dtMirOps Generated	The number of mirror operations (create, modify, or delete) that have been transmitted by the mirroring process
	dtMirBytes Generated	The number of bytes that have been transmitted by the mirroring process
	dtRepOps Generated	The number of operations (create, modify, or delete) that have been transmitted by the replication process
	dtRepBytes Generated	The number of bytes that have been transmitted by the replication process
	dtFailedMirror Count	The number of operations that failed to mirror because they couldn't be read on the source
	dtFailedRepCount	The number of operations that failed to be replicated because they could not be read on the source
	dtActFailCount	The number of activation code errors
General	**dtAutoDisCount**	The number of auto-disconnects
	dtAutoReCount	The number of auto-reconnects
	dtDriverQueue Percent	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
	dtRetryCount	The number of file operations that have been retried
	dtOpsDropped Count	The number of file operations that have failed and will not be retried
Security	**dtLoginCount**	The number of successful logins
	dtFailedLogin Count	The number of unsuccessful logins
Connection	**dtConnection Count**	The number of active connections between machines
	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.
	dtconConnect Time	The duration of time since the connection was first established

Object Type	Statistic	Description
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).
	dtconOpsInCmd Queue	The number of operations (create, modify, or delete) in the retransmit queue on the source
	dtconOpsInAck Queue	The number of operations (create, modify, or delete) waiting for verification acknowledgements from the target
	dtconOpsInRep Queue	The number of replication operations (create, modify, or delete) in the queue
	dtconOpsInMir Queue	The number of mirror operations (create, modify, or delete) in the queue
	dtconBytesInRep Queue	The number of bytes in the replication queue
	dtconBytesInMir Queue	The number of bytes in the mirror queue
	dtconOpsTx	The total number of operations (create, modify, or delete) transmitted to the target
	dtconBytesTx	The total number of bytes transmitted to the target
	dtconBytes CompressedTx	The total number of compressed bytes transmitted to the target
	dtconOpsRx	The total number of operations (create, modify, or delete) received from the target
	dtconBytesRx	The total number of bytes received from the target
	dtconResentOp Count	The number of operations that were resent because of acknowledgement errors

SNMP traps

The following tables identify the Double-Take traps provided through SNMP.

Object Type	Trap	Description
Kernel	**dttrapKernel Started**	Double-Take has started
	dttrapKernel Stopped	Double-Take has stopped
License	**dttrapLicense ViolationStarting Source**	The source cannot be started due to a license violation
	dttrapLicense ViolationOn Network	A Double-Take serial number conflict was identified on the network
Source	**dttrapSource Started**	Double-Take source component has started
	dttrapSource Stopped	Double-Take source component has stopped
Target	**dttrapTarget Started**	Double-Take target component has started
	dttrapTarget Stopped	Double-Take target component has stopped
Connection	**dttrapConnection Requested**	The source has requested a connection to the target
	dttrapConnection RequestReceived	The target has received a connection request from the source
	dttrapConnection Succeeded	The source to target connection has been established
	dttrapConnection Pause	The source to target connection has paused
	dttrapConnection Resume	The source to target connection has resumed
	dttrapConnection Failed	The source to target connection was not successful
	dttrapConnection Lost	The source to target connection has been disconnected
	dttrapMemory LimitReached	The Double-Take memory pool limit has been reached
	dttrapMemory LimitRemedied	The memory pool usage is below the maximum limit specified
	dttrapAuto Reconnect	Auto-reconnect needs to make a new connection
	dttrapScheduled ConnectStart	A scheduled connection has been established
	dttrapScheduled ConnectEnd	A scheduled end connection has been reached and the connection has been disconnected
	dttrapAuto DisconnectWrite Queue	Auto-disconnect has forced the queue to be written to disk

Object Type	Trap	Description
Connection	**dttrapAuto DisconnectPause Transmission**	Auto-disconnect requested that the source pause any operation (create, modify, or delete) sending
	dttrapAuto DisconnectEnd Connection	Auto-disconnect has intentionally dropped the connection
	dttrapAuto Disconnect Shutdown	Auto-disconnect forced Double-Take to shutdown
Replication	**dttrapReplication Start**	Replication has started
	dttrapReplication Stop	Replication has stopped
Mirroring	**dttrapMirrorStart**	Mirroring has started
	dttrapMirrorStop	Mirroring has stopped
	dttrapMirror Pause	Mirroring has paused
	dttrapMirror Resume	Mirroring has resumed
	dttrapMirrorEnd	Mirroring has ended
Verification	**dttrapVerification Start**	Verification has started
	dttrapVerification End	Verification has ended
	dttrapVerification Failure	Verification has failed
Restoration	**dttrapRestore Started**	Restoration has started
	dttrapRestore Complete	Restoration is complete
Replication Sets	**dttrapRepSet Modified**	Replication has been modified
Failover	**dttrapFailover ConditionMet**	Manual intervention is required because failover has detected a failed source machine
	dttrapFailoverIn Progress	Failover is occurring
	dttrapTargetFull	The target is full

Server Settings

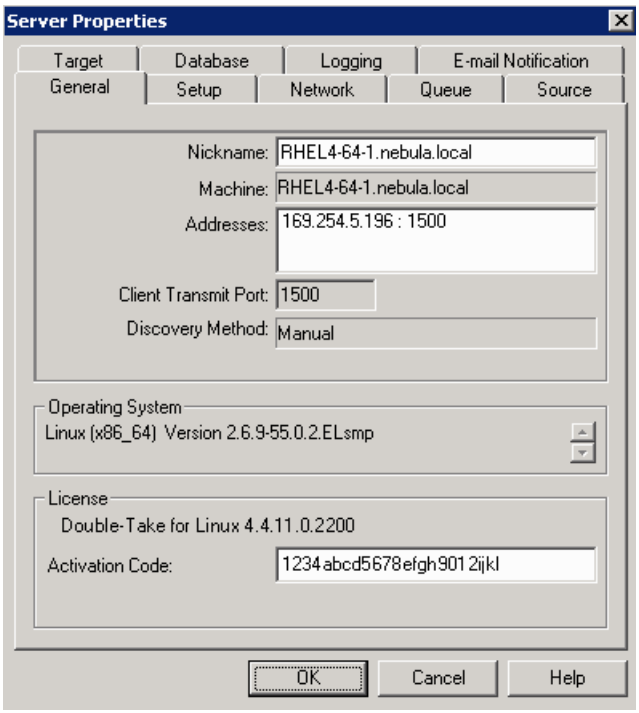
16

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

- ***General*** on page 16-2
- ***Setup*** on page 16-4
- ***Network*** on page 16-6
- ***Queue*** on page 16-7
- ***Source*** on page 16-9
- ***Target*** on page 16-11
- ***Database*** on page 16-12
- ***Logging*** on page 16-13
- ***E-mail notification*** on page 16-14

General

The **General** tab identifies the server.



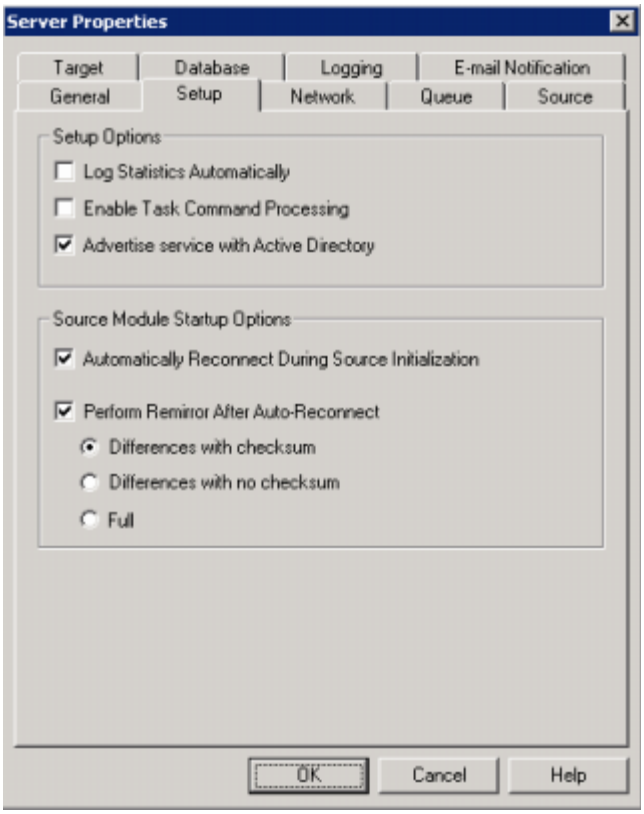
- **Nickname**—A nickname is saved in the Management Console workspace, therefore, it only appears in the Management Console on this server. It is not communicated across the network. If you export a workspace and use it on another Double-Take server, the server nickname will appear there also.
- **Machine**—This 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 Management Console uses to send commands to a server. This port cannot be modified.
- **Discovery Method**—This field indicates the method in which the Management Console identifies the Double-Take server.
  - **Manual**—A Double-Take server was manually inserted into the Management Console server tree.
  - **Broadcast Heartbeat**—A Double-Take server is broadcasting Double-Take heartbeats.
- **Operating System**—The server's operating system version is displayed.
- **Double-Take Version Information**—The Double-Take version number and build number are displayed.

- **Activation Code**—This is the Double-Take license which is required on every Double-Take machine. There are three licenses available.
  - Evaluation—A license that has an end 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 license that is available on a per machine basis. Each machine is required to have a unique Double-Take license whether it is functioning as a source, target, or both. A single license can only be used on one machine on a network.
  - Site—A license that is available to register every machine with the same Double-Take license. This license is designed to be used on multiple machines on a network.

  The activation code is a 16 character, alpha-numeric code. This code determines the Double-Take license that is in use and can be changed without reinstalling Double-Take if your Double-Take license changes.

Setup

The **Setup** tab contains various setup options.



- **Log Statistics Automatically**—If enabled, Double-Take statistics logging will start automatically when Double-Take is started.
- **Enable Task Command Processing**—Task command processing is a Double-Take 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 Management Console, but it can only be initiated through the scripting lanugage. See the *Scripting Guide* for more information.

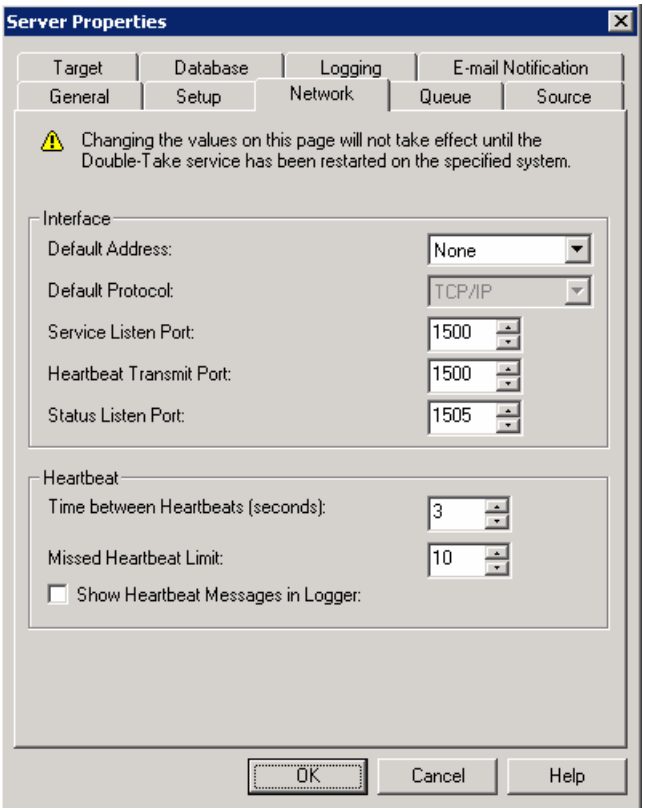
  If you disable this option on a source server, you can still submit tasks to be processed on a target, although task command processing must be enabled on the target.
- **Advertise service with Active Directory—**This option is for Windows only.
- **Automatically Reconnect During Source Initialization**—If enabled, Double-Take will automatically reconnect any connections that it automatically disconnected.
- **Perform Remirror After Auto-reconnect**—If enabled, Double-Take 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.

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

Network

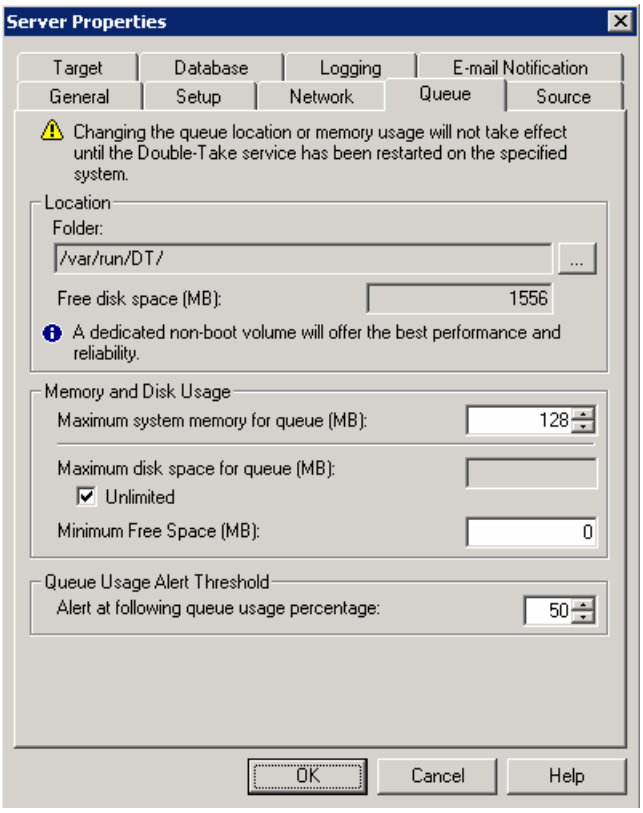
The **Network** tab contains settings that control network communications. Any changes made to this tab will not take effect until the Double-Take daemon has been restarted on the server.



- **Default Address**—On a machine with multiple NICs, you can specify which address Double-Take 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 communications is the TCP/IP protocol. In the future, Double-Take may support other communication protocols.
- **Service Listen Port**—Double-Take servers use the **Service Listen Port** to send and receive commands and operations between two Double-Take servers.
- **Heartbeat Transmit Port**—A Double-Take server sends its heartbeats to the **Heartbeat Transmit Port**.
- **Status Listen Port**—Double-Take servers use the **Status Listen Port** to listen for requests from the Management Console and other clients.
- **Time Between Heartbeats**—All Double-Take servers transmit a heartbeat. This heartbeat allows other Double-Take servers and Double-Take clients to locate and identify the Double-Take servers. The heartbeat is a broadcast UDP transmission. This heartbeat can be disabled, but if it is, Double-Take will not auto-detect the Double-Take servers to populate the Management 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 log. Enabling this option will cause your logs to fill up faster.

Queue

The **Queue** tab contains settings for Double-Take queue usage.



- **Folder**—This is the location where the disk queue will be stored. Double-Take 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.

  NOTE: Scanning the Double-Take 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 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 will use less system memory, but you will queue to disk sooner which may impact system performance. If you set it higher, Double-Take 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 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 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, you might configure the application to use 512 MB and Double-Take 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.

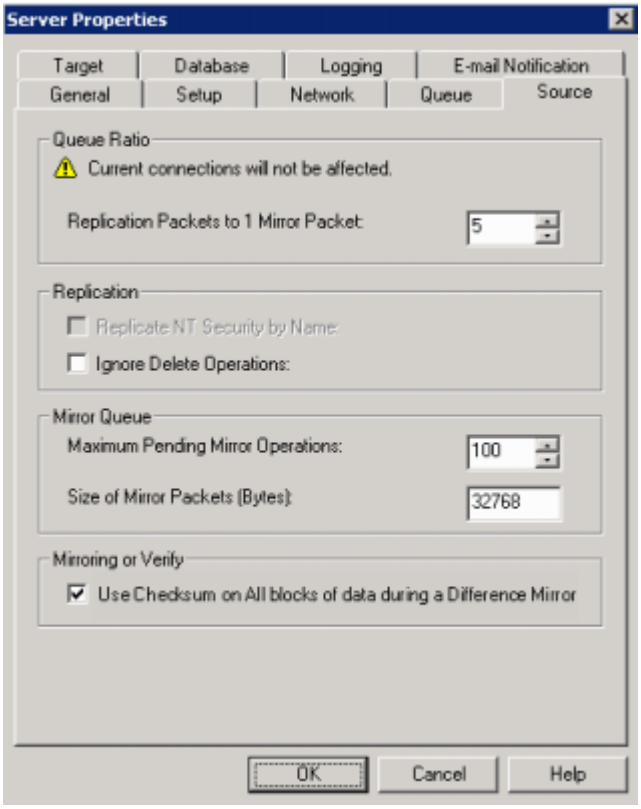
- **Maximum disk space for queue**—This is the maximum amount of disk space, in MB, in the specified **Folder** that can be used for Double-Take disk queuing, or you can select **Unlimited** which will allow the queue usage to automatically expand whenever the available disk space expands. When the disk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

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

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

Source

The **Source** tab contains settings specific to the source module of Double-Take.



- **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.

NOTE:	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.

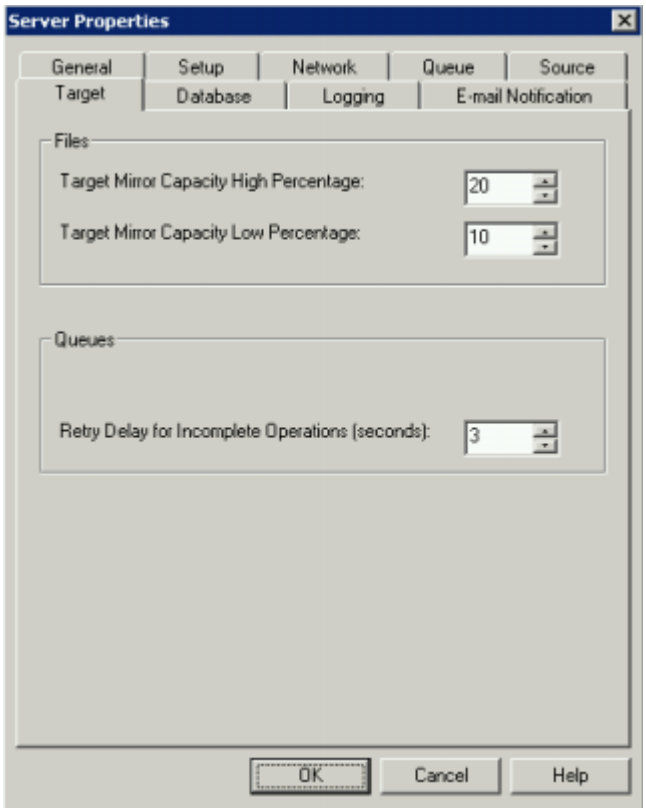
- **Maximum Pending Mirror Operations**—This option is the maximum number of mirror operations that are queued on the source. The default setting is `1000`. If, during mirroring, the mirror queued statistic regularly shows low numbers, for example, less than `50`, this value can be increased to allow Double-Take to queue more data for transfer.
- **Size of Mirror Packets**—This option determines the size of the mirror packets that Double-Take 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 will assume files are synchronized if their attributes match.

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

Target

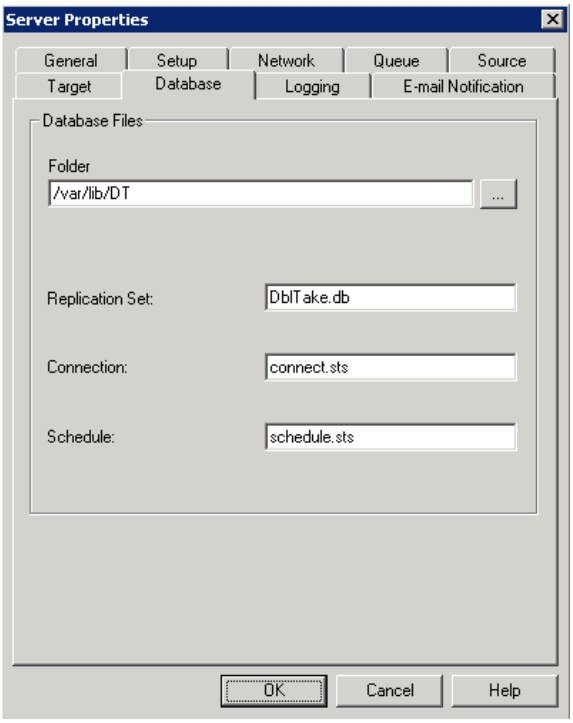
The **Target** tab contains settings specific to the target module of Double-Take.



- **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`.

Database

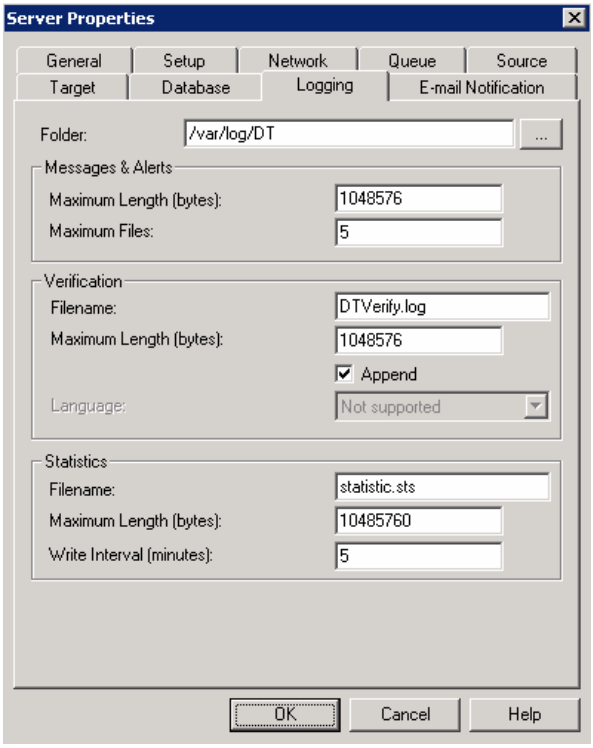
The **Database** tab identifies the Double-Take database files.



- **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 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.

Logging

The **Logging** tab contains settings for various Double-Take log files.



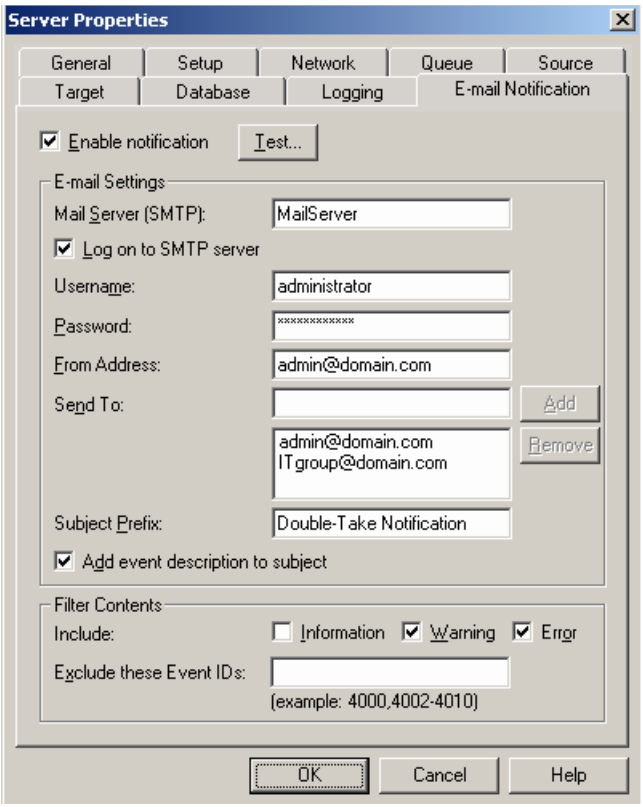
- **Folder**—Specify the directory where each of the log files on this tab are stored. The default location is the directory where the Double-Take program files are installed.
- **Messages & Alerts**
  - **Maximum Length**—Specify the maximum length of the client and daemon log files. The default size is 1048576 bytes and is limited by the available hard drive space.
  - **Maximum Files**—Specify the maximum number of Double-Take 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 begins overwriting the oldest data in the file.
  - **Write Interval**—Specify how often Double-Take writes to the statistics log file. The default is every 5 minutes.

E-mail notification

The **E-mail Notification** tab allows you to specify e-mail notification settings. When enabled, the messages logged to the Linux System Log will be sent as e-mail messages to the addresses specified. 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.

In order to enable e-mail notification, the server must have Internet Explorer 5.0 or later.

Select or deselect **Enable notification** to turn this feature on or off. Any notification settings specified are retained when **Enable notification** is disabled.



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

NOTE:	Specifying an SMTP server is the preferred method because it provides a direct connection between the mail server and Double-Take, 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 will attempt to use the Linux mail command. The success will depend on how the local mail system is configured. Double-Take 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 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 e-mail message. The address is limited to 256 characters.

- **Send To**—Specify the e-mail address that each Double-Take 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 e-mail message. This will help distinguish Double-Take messages from other messages. This field is optional.

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

NOTE: You can test e-mail notification by specifying the options on the E-mail Notification tab and clicking **Test**. 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 gete-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 e-mail messages are not blocked.

Security 17

To ensure protection of your data, Double-Take offers multi-level security using native operating system security features. Privileges are granted through membership in user groups defined on each machine running Double-Take. To gain access to a particular Double-Take 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 security groups. Once a valid user name and password have been provided and the Double-Take source or target has verified membership in one of the Double-Take 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 is granted on one of the following three levels:

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

Security access levels

The following table identifies which key Double-Take features are available depending on the security access granted.

Double-Take Feature	Administrator Access	Monitor Access
Modify Replication Sets	✓	
View Replication Sets and Rules	✓	✓
Control Connections, Mirroring, Replication, Verification, Restoration, Failover	✓	
View Connection, Mirroring, Replication, Verification, Restoration Processing Statistics	✓	✓
View Double-Take Program Settings	✓	✓
Modify Double-Take Program Settings	✓	

Security advantages and considerations

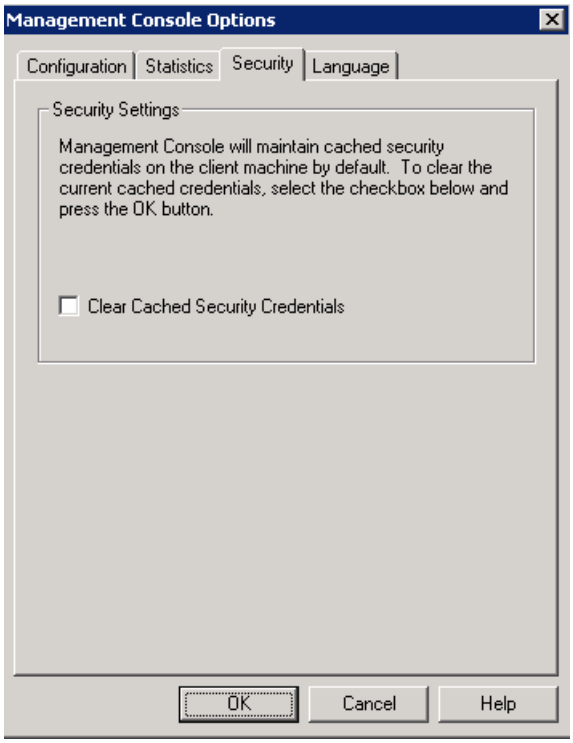
Double-Take security provides machine-based protection allowing the network administrator to specify the individuals that can access all of the Double-Take features as well as those that only have access to the Double-Take statistics. This security prevents unauthorized users from modifying critical Double-Take configurations like the data included or excluded from a replication set, changing a one-to-one configuration to a one-to-many configuration by adding another connection, or initiating a mirror or stopping replication.

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

Clearing maintained security credentials

To remove cached credentials, follow these steps.

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**.

Index

A

Active Directory
alternate data streams 4-3

B

block target paths 16-11

C

centralized backup 1-1
chained configuration 1-11
client logs 16-13
client overview 3-1
compression 9-13
configurations 1-7
connection
  automatic reconnections 9-12
  Connection Manager 6-1
  Connection Wizard 5-1
  disconnecting 9-13
  mirroring 6-3
  monitoring 9-2
  overview 9-1
  simulating 8-1
Connection Manager 6-1
Connection Wizard 5-1
credentials
  Failover Control Center 3-4
  Management Console 17-2

D

database settings 16-12
disaster recovery
  overview 1-1
DTInfo 8-1
DTStat 15-19
dynamic volumes 4-3

E

e-mail notification 16-14
encrypted files 4-3
event messages 15-9

F

failback 13-6
failover
  configuring 13-1
  data 13-5
  editing configuration 13-3
  failback 13-6
  manual intervention 13-5
  monitoring 13-4
  overview 1-5, 13-1
  removing 13-4
  testing 13-5
Failover Control Center
  credentials 3-4
  overview 3-4
  ports 3-4
  refresh rate 3-4
failure monitoring 1-5
firewall 7-1

H

hard links 4-3
heartbeats 16-6
high availability
  overview 1-1

I

ignore delete operations 16-9
in-band control
  See task command processing
insert servers Management Console 7-2
installation
  system requirements 2-1
intervention 13-5

L

logging
  clients 15-2
  daemon 15-2
  messages 15-9
  settings 16-13
  SNMP 15-26
  system log 15-9
  verification 11-3
logging on/off 3-2

M

Management Console
  credentials 17-2
  insert servers 7-2
  message window 9-6
  overview 3-1
  workspaces 3-3
manual intervention 13-5
many-to-one configuration 1-9
message window 9-6
mirroring
  automatic 10-3
  connection 6-3
  manual controls 10-1
  options compared 10-2, 10-4
  overview 1-3
  remirroring options compared 10-4
monitoring
  connections 9-2
  failover 13-4
  message window 9-6
  statistics 15-19
  tools 15-1
mount points 4-3

N

NAT 7-1
network settings 16-6

O

one-to-many configuration 1-10
one-to-one configuration 1-7, 1-8
orphan files 10-5
overview
  operations 1-3
  product 1-1

P

ports
  Failover Control Center 3-4
  overview 7-1
  server 16-6

Q

queuing 9-8, 9-10, 16-7

R

refresh rate 3-4
remirroring
  options compared 10-4
reparse points 4-3
replication
  capabilities 4-3
  overview 1-4
  starting 10-6
replication set
  copying 4-7
  creating 4-5, 4-6
  deleting 4-7
  modifying 4-7
  renaming 4-7
  replication capabilities 4-3
  size 4-8
resources 1-14
restoration 14-1
  overview 1-6

S

security
  credentials
    Failover Control Center 3-4
    Management Console 17-2
  overview 17-1
security by name 16-9
server properties
  database 16-12
  e-mail notification 16-14
  general 16-2
  logging 16-13
  network 16-6
  queue 16-7
  setup 16-4

source 16-9
target 16-11
setup settings 16-4
simulating a connection 8-1
single server configuration 1-12
SMTP server 16-14
SNMP 15-26
source settings 16-9
sparse files 4-3
statistics
DTStat 15-19
log 16-13
SNMP 15-26
system information 8-1
system log 15-9
system requirements 2-1

T

target controls 12-1
target settings 16-11
task command processing 16-4
TDU 8-1
Throughput Diagnostics Utility 8-1
transmission
controls 12-1
manual controls 12-1
overview 12-1
scheduling 12-2

V

verification
log 11-3, 16-13
manual control 11-1
overview 11-1
scheduling 11-2
virtual server configuration 1-13
virus protection 4-2

W

workspaces 3-3