$\textbf{Double-Take}^{\mathbb{R}} \, \mathsf{AVAILABILITY}^{^{\mathsf{TM}}}$

Version 7.0 Double-Take Availability for Linux User's Guide





Notices

Double-Take Availability for Linux User's Guide Version 7.0, Friday, December 13, 2013

Check the Vision Solutions support web site at http://www.VisionSolutions.com/SupportCentral for the most up-to-date version of this documentation.

- Product Updates—Check your service agreement to determine which updates and new releases
 you may be eligible for. Product updates can be obtained from the support web site at
 http://www.VisionSolutions.com/SupportCentral.
- Sales—If you need maintenance renewal, an upgrade activation code, or other sales assistance, contact your reseller/distributor or a Vision Solutions sales representative. Contact information is available on the Vision Solutions Worldwide Locations and Contacts web page at http://www.VisionSolutions.com/Company/Vision-HA-Locations.aspx.
- Technical Support—If you need technical assistance, you can contact CustomerCare. All basic
 configurations outlined in the online documentation will be supported through CustomerCare. 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
 CustomerCare and they will direct you to the correct service provider. To contact CustomerCare, you
 will need your serial number and activation code. Contact information is available on the Vision
 Solutions CustomerCare web page at http://www.VisionSolutions.com/Support/SupportOverview.aspx.
- Professional Services—Assistance and support for advanced configurations may be referred to a Pre-Sales Systems Engineer or to Professional Services. For more information, see the Windows and Linux tab on the Vision Solutions Consulting Services web page at http://www.VisionSolutions.com/Services/Consulting-Services.aspx.
- **Training**—Classroom and computer-based training are available. For more information, see the Double-Take Product Training web page at http://www.VisionSolutions.com/Services/DT-Education.aspx.
- Documentation—Please forward any comments or suggestions about this documentation to documentation-Double-Take@VisionSolutions.com.
- Linux man pages—Man pages are installed and available on Double-Take Linux servers. These documents are bound by the same Vision Solutions license agreement as the software installation.

This documentation is subject to the following: (1) Change without notice; (2) Furnished pursuant to a license agreement; (3) Proprietary to the respective owner; (4) Not to be copied or reproduced unless authorized pursuant to the license agreement; (5) Provided without any expressed or implied warranties, (6) Does not entitle Licensee, End User or any other party to the source code or source code documentation of anything within the documentation or otherwise provided that is proprietary to Vision Solutions, Inc.; and (7) All Open Source and Third-Party Components ("OSTPC") are provided "AS IS" pursuant to that OSTPC's license agreement and disclaimers of warranties and liability.

Vision Solutions, Inc. and/or its affiliates and subsidiaries in the United States and/or other countries own/hold rights to certain trademarks, registered trademarks, and logos. Hyper-V and Windows are registered trademarks of Microsoft Corporation in the United States and/or other countries. Linux is a registered trademark of Linus Torvalds. vSphere is a registered trademark of VMware. All other trademarks are the property of their respective companies. For a complete list of trademarks registered to other companies, please visit that company's website.

© 2013 Vision Solutions, Inc. All rights reserved.

Double-Take Availability overview	7
Core operations	
Supported configurations	
Selecting a protection type	
Volume 1 Linux files and folders protection	19
Chapter 1 Files and folders requirements	
Chapter 2 Installation	
Installation and upgrade notes	
Installing or upgrading Double-Take Linux servers	
Installing or upgrading Double-Take Windows client	
Removing (erasing) Double-Take	29
Chapter 3 DTSetup	
Running DTSetup	
Setup tasks	
Activating your server	
Modifying security groups	
Replication configuration	
Configuring file system replication	
Configuring block device replication	36
Configuring server settings	38
Configuring driver performance settings	39
Starting and stopping the daemon	40
Starting DTCL	41
Viewing documentation and troubleshooting tools	
DTSetup menus	
Chapter 4 Clients	
Replication Console	45
Using Replication Console workspaces	46
Clearing stored security credentials	
Failover Control Center	
Setting the frequency of Failover Control Center console refreshes	49
Chapter 5 Data protection	
Establishing a data connection using the automated Connection Wizard	
Creating a replication set	
Establishing a connection manually using the Connection Manager	56
Establishing a connection across a NAT or firewall	
Simulating a connection	
Chapter 6 Protection monitoring	
Monitoring a data workload	
Log files	
Viewing the log files through a text editor	
Viewing the Double-Take log file through the Replication Console	
Configuring the properties of the Double-Take log file	
Double-Take log messages	
Monitoring the Linux system log	
E-mailing system messages	92

Statistics	95
Configuring the properties of the statistics file	96
Viewing the statistics file	97
Statistics	99
SNMP	105
Configuring SNMP on your server	106
SNMP traps	
SNMP statistics	
Chapter 7 Connections	
Data queues	
Queuing data	
Auto-disconnect and auto-reconnect	
Reconnecting automatically	121
Pausing and resuming target processing	
Disconnecting a connection	
Chapter 8 Mirroring	
Stopping, starting, pausing, or resuming mirroring	
Mirroring automatically	
Removing orphan files	
Chapter 9 Replication	131
Replication capabilities	
Replication sets	134
Creating a replication set	136
Creating or modifying replication rules manually	139
Selecting a block device for replication	
Modifying a replication set	
Renaming and copying a replication set	143
Calculating replication set size	
Exporting and importing a replication set	146
Deleting a replication set	147
Starting replication	148
Inserting tasks during replication	149
Chapter 10 Verification	
Verifying manually	151
Verifying on a schedule	
Configuring the verification log	154
Chapter 11 Data transmission	
Stopping, starting, pausing, and resuming transmission	157
Scheduling data transmission	157
Limiting transmission bandwidth	162
Compressing data for transmission	164
Chapter 12 Failover	166
Configuring failover monitoring	167
WAN considerations	
Protecting NFS exports	172
Protecting Samba shares	
Editing failover monitoring configuration	
Monitoring failover monitoring	175
Failing over	

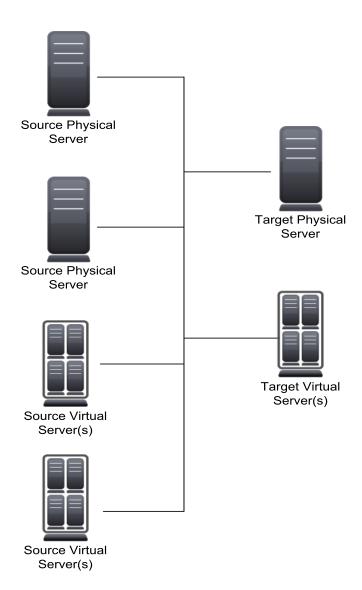
	lover monitoring configuration	
	pack and restoration	
Restoring the	n failing back	180
Failing back th	nen restoring	184
Chapter 14 Serv	er settings	187
Identifying a s	erver	188
Licensing a se	erver	190
Configuring so	erver startup options	193
Configuring n	etwork communication properties for a server	195
Queuing data		197
Configuring so	ource data processing options	200
	arget data processing options	
	e Double-Take database storage files	
	names for logging and statistics	
	tem messages	
	ırity	
	nd off	
	o ESX appliance protection	
	ng started	
	ESX appliance requirements	
	Double-Take virtual recovery appliance	
	Double-Take Console	
	al recovery appliance to the console	
	Double-Take virtual recovery appliance	
	ble-Take Availability on Linux	
	g the Double-Take Console	
	x source to the console	
	I server to ESX appliance job	
	ging and controlling full server to ESX appliance jobs	
	erver to ESX appliance job details	
	ıll server to ESX appliance job	
	erver to ESX appliance job	
	server to ESX appliance job log	
	e statistics file	
	g over full server to ESX appliance jobs	
	le-Take Console	
	Console requirements	
	ons	
	og files through the Double-Take Console	
	ging servers	
Providing serv	ver credentials	287
Viewing serve	er details	288
	properties	
	nsing	
	er logs	
License mana	agement and activation	295
Managing tl	he Double-Take license inventory	296

Licensing a server or appliance	
Activating a single license	
Activating multiple licenses	
Managing VMware servers	

Double-Take Availability overview

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

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



Core operations

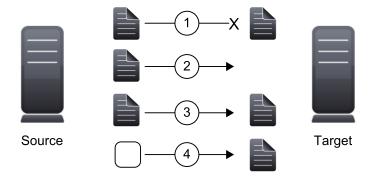
Double-Take performs three basic types of operations.

- See Mirroring on page 8—The initial copy or subsequent resynchronization of selected data
- See Replication on page 9—The on-going capture of byte-level file changes
- See Failover on page 10—The ability to stand-in for a server, in the event of a failure

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

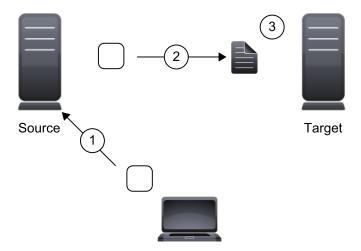


- 1. Identical files are not mirrored.
- 2. New files are mirrored.
- 3. Different files can be mirrored.
- 4. Checksums can calculate blocks of data to be mirrored.

Replication

Replication is the real-time transmission of file changes. Unlike other related technologies, which are based on a disk driver or a specific application, the Double-Take 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.



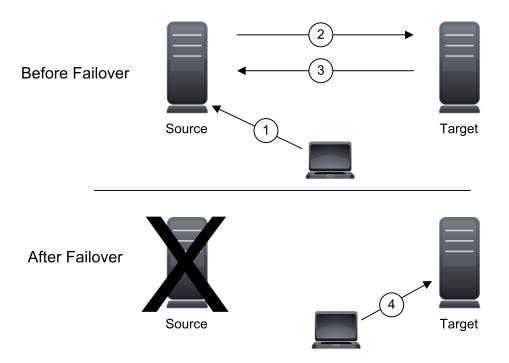
- 1. A user or application updates part of a file.
- 2. Only the changed portion of the file is replicated to the target.
- 3. An up-to-date copy of the file is maintained on the target.

Failover

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

Double-Take monitors the source status by tracking requests and responses exchanged between the source and target. When a monitored source does not respond to the target's 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 identity of the failed source, and 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.



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

Supported configurations

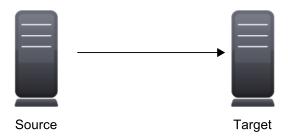
Double-Take is an exceptionally flexible product that can be used in a wide variety of network configurations. To implement Double-Take effectively, it is important to understand the possible configuration options and their relative benefits. Double-Take configurations can be used independently or in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



Not all types of jobs support all of these configurations. See the requirements of each job type to determine which configurations are supported.

- See One to one, active/standby on page 12
- See One to one, active/active on page 13
- See Many to one on page 14
- See One to many on page 15
- See Chained on page 16
- See Single server on page 17

One to one, active/standby



Description

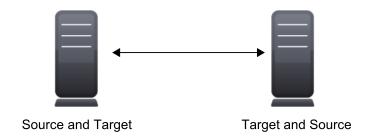
One target server, having no production activity, is dedicated to support one source server. The source is the only server actively replicating data.

Applications

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

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

One to one, active/active



Description

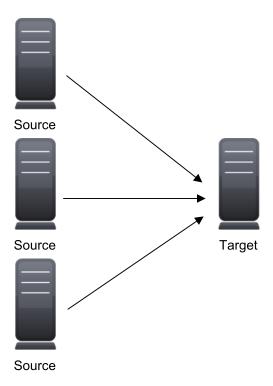
Each server acts as both a source and target actively replicating data to each other

Applications

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

- 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 protected data set. Any overlap will cause an infinite loop.
- To support the production activities of both servers during failover without reducing performance, each server should have sufficient disk space and processing resources.
- Failover and failback scripts must be implemented to avoid conflict with the existing production applications.
- You must pause the target when backing up database files on the target.

Many to one



Description

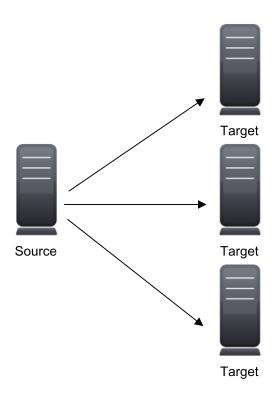
Many source servers are protected by one target server.

Applications

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

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

One to many



Description

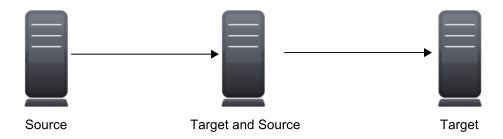
One source server sends data to multiple target servers. The target servers may or may not be accessible by one another.

Applications

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

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

Chained



Description

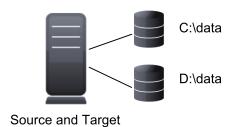
The source servers sends replicated data to a target server, which acts as a source server and sends data to a final target server, which is often offsite.

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.

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

Single server



Description

Source and target 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.

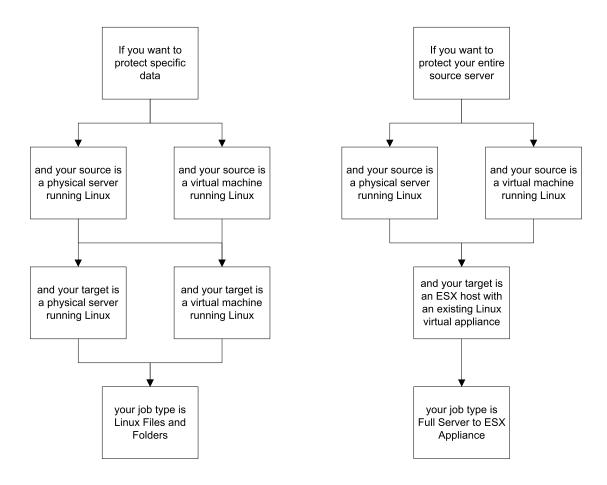
Applications

This configuration is useful upgrading storage hardware while leaving an application online. Once the data is mirrored, you can swap the drive in the disk manager. If the source and target copies of the data are located on different drives, this configuration supports high availability of the data in the event that the source hard drive fails.

- 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.

Selecting a protection type

Double-Take is an exceptionally flexible product that can be used in a wide variety of network configurations. However, this flexibility can make it difficult to determine which Double-Take job is right for your environment. Knowing what you want to protect on your source is the key to determine which Double-Take job to use. Review the following decision trees to find which job type is best for your needs and environment.



Volume 1 Linux files and folders protection

Create a files and folders job when you want to protect data. You can also use it to protect applications, such as Oracle or MySQL, however you will need to use your own customized failover and failback scripts to start and stop services during failover and failback. This job type does not protect a server's system state. Linux files and folders protection is distinct from full server to ESX appliance protection. For example, it has its own requirements and client components. Use the following links to access information and steps specific to files and folder protection.

- Files and folders requirements on page 20
- Installation on page 24
- DTSetup on page 30
- Clients on page 44
- Data protection on page 50
- Protection monitoring on page 63
- Connections on page 113
- Mirroring on page 124
- Replication on page 131
- Verification on page 150
- Data transmission on page 156
- Failover on page 166
- Failback and restoration on page 179
- Server settings on page 187
- Security on page 209

Chapter 1 Files and folders requirements

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

- Source and target server requirements on page 20
- Console requirements on page 23

Source and target server requirements

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

Red Hat Enterprise Linux and CentOS

Operating system version—4.8 through 4.9

Kernel type for x86 (32-bit) architectures—Default, SMP, HugeMem

Kernel type for x86 64 (64-bit) architectures—Default, SMP, LargeSMP

File system—Ext3

Red Hat Enterprise Linux, CentOS, and Oracle Enterprise Linux

Operating system version—5.8 through 5.9

Kernel type for x86 (32-bit) architectures—Default, SMP, Xen, PAE

Kernel type for x86-64 (64-bit) architectures—Default, SMP, Xen

File system—Ext3, Ext4, XFS

Red Hat Enterprise Linux, CentOS, and Oracle Enterprise Linux

Operating system version—6.3 through 6.4

Kernel type for x86 (32-bit) architectures—Default

Kernel type for x86-64 (64-bit) architectures—Default

File system—Ext3, Ext4, XFS (64-bit only)

SUSE Linux Enterprise

Operating system version—10.3 and 10.4

Kernel type for x86 (32-bit) architectures—Default, SMP, BigSMP, Xen, XenPAE

Kernel type for x86-64 (64-bit) architectures—Default, SMP, Xen

File system—Ext3, ReiserFS, XFS

SUSE Linux Enterprise

Operating system version—11.1 through 11.2

Kernel type for x86 (32-bit) architectures—Default, Xen, XenPAE, VMI

Kernel type for x86-64 (64-bit) architectures—Default, Xen

File system—Ext3, ReiserFS, XFS

Ubuntu

Version—10.04.3

Kernel version—2.6.32-33

Kernel type for x86 (32-bit) architectures—386, Generic, Generic-pae

Kernel type for x86-64 (64-bit) architectures—Generic, Preempt, Server

File system—Ext2, Ext3, Ext4, XFS

Ubuntu

Version—10.04.4

Kernel version—2.6.32-38

Kernel type for x86 (32-bit) architectures—386, Generic, Generic-pae

Kernel type for x86-64 (64-bit) architectures—Generic, Preempt, Server

File system—Ext2, Ext3, Ext4, XFS

Ubuntu

Version—12.04.0

Kernel version—3.2.0-23

Kernel type for x86 (32-bit) architectures—Generic, Virtual, Generic-pae,

Low latency

Kernel type for x86-64 (64-bit) architectures—Generic, Virtual, Low latency

File system—Ext2, Ext3, Ext4, XFS

Ubuntu

Version—12.04.1

Kernel version—3.2.0-29

Kernel type for x86 (32-bit) architectures—Generic, Generic-pae, Virtual

Kernel type for x86-64 (64-bit) architectures—Generic, Virtual

File system—Ext2, Ext3, Ext4, XFS

Ubuntu

Version—12.04.2

Kernel version—3.5.0-23

Kernel type for x86 (32-bit) architectures—Generic

Kernel type for x86-64 (64-bit) architectures—Generic

File system—Ext2, Ext3, Ext4, XFS

Ubuntu

Version—12.04.3

Kernel version—3.8.0-29

Kernel type for x86 (32-bit) architectures—Generic

Kernel type for x86-64 (64-bit) architectures—Generic

File system—Ext2, Ext3, Ext4, XFS



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

Oracle Enterprise Linux support is for the mainline kernel only, not the Unbreakable kernel.

Double-Take does not support stacking filesystems, like eCryptFS.

You must have sshd (or the package that installs sshd), lsb, parted, /usr/sbin/dmidecode, and /usr/bin/which on your Linux servers before you can install and use Double-Take Availability. See your operating system documentation for details on these packages and utilities.

- **System Memory**—The minimum system memory on each server should be 1 GB. The recommended amount for each server is 2 GB.
- **Disk Usage**—The amount of disk space required for the Double-Take program files is approximately 85 MB. About 45 MB will be located on your /(root) partition, and the remainder will be on your /usr partition. You will need to verify that you have additional disk space for Double-Take queuing, logging, and so on. Additionally, on a target server, you need sufficient disk space to store the replicated data from all connected sources, allowing additional space for growth.
- **Protocols**—Your servers must have TCP/IP. IPv4 is the only supported version.
- **Ports**—Port 1501 is used for localhost communication. Ports 1500, 1505, 1506, 6325, and 6326 are used for component communication and must be opened on any firewall that might be in use.
- IP address and subnet configuration—Because of limitations in the way the Linux kernel handles IP address aliases, do not mix subnets on the eth0 network interface. Failover should not cause problems in this configuration, but you will lose IP addresses during failback. Therefore, if you must mix subnets on a single interface, use eth1 or higher.
- Name resolution—Your servers must have name resolution or DNS. The Replication Console
 and interactive text client (DTCL -i) will fail if there is no DNS entryor way for a server to resolve
 server names. For details on name resolution options, see your Linux documentation or online
 Linux resources.

- **Security**—Double-Take security is granted through membership in user groups. The groups can be local or LDAP (Lightweight Directory Access Protocol). A user must provide a valid local account that is a member of the Double-Take security groups.
- **SELinux policy**—SELinux should be disabled on the source and target.
- VMware Tools—Any VMWare guest running Double-Take should have the appropriate VMWare Tools package installed.
- Hard links—If you have hard links outside of the data set you are protecting, and they link to files
 inside the data set you are protecting, Double-Take will not mirror or replicate the hard links which
 could lead to differences on the target.

Console requirements

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

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

Chapter 2 Installation

Review *Files and folders requirements* on page 20 and *Installation and upgrade notes* on page 25 then use the appropriate instructions from the following list to meet your goal.

- Installing or upgrading Double-Take Linux servers on page 26—Use these instructions if you are installing or upgrading on a Linux server.
- Installing or upgrading Double-Take Windows client on page 28—Use these instructions if you want to install or upgrade on a Windows client.

Installation and upgrade notes

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

- Because Double-Take 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.
- A script called DTInfo.sh is included in the installation. This script can be run via DTSetup to
 collect configuration data for use when reporting problems to technical support. Depending on the
 type of diagnostic information gathering you select, DTInfo can gather Double-Take log files,
 Double-Take and system settings, network configuration information, and other data which may
 be necessary for technical support to troubleshoot issues. After running the script, a .tar.gz is
 automatically created with the information gathered. You must have root (or uid 0 equivalent) to
 execute the diagnostics or to copy or read the resulting file.
- Double-Take 7.0 is interoperable back to version 4.7 but is restricted to the following limitations.
 The Double-Take clients can only control the same or older releases. To accommodate rolling upgrades, older sources can connect to newer targets, but newer sources cannot connect to older targets.
 - 4.7 client—Supports 4.7 source and target, but does not support 6.0 or 7.0 source or target
 - **6.0 client**—Supports 4.7 or 6.0 source and target as long as the target is the same or newer than the source, but does not support 7.0 source or target
 - 7.0 client—Supports 4.7, 6.0, or 7.0 source and target as long as the target is the same or newer than the source
- When performing a rolling upgrade, update the target servers first. After the upgrade is complete
 and the target daemon is restarted, the sources will automatically reconnect to the targets.
 Upgrade the sources when convenient.
- If you are using a chained configuration, update the last target first, then update the middle server acting as both a source and target, and update the production source last.
- If you are using a configuration where the source is an older version than the target, you will not be
 able to restore from the newer version target back to the older version source. You must upgrade
 the source to the same version as the target before restoring.

Installing or upgrading Double-Take Linux servers

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



The instructions below are formatted for Red Hat and similar operating systems. You may need to modify the commands to fit your operating system. For example, on Ubuntu you should use dpkg -i to install the .deb file. See your operating system reference documentation for appropriate installation commands for your server.

If you are using SLES version 10.x, you may get an installation database warning regarding an exclusive database lock. You can safely disregard this message.

- 1. Determine the installation package that is appropriate for your operating system. There are separate .rpm installation packages for 32-bit and 64-bit architectures. Make sure you are installing the correct .rpm file. If you are uncertain about the architecture of your machine, you can use the uname -m command to determine it. Additionally, the version_numbers in the installation file name will vary and will correspond to the version of Double-Take you are installing. For example, if you are installing version 7.0.0.1124.0, the installation files would be DoubleTake-7.0.0.1124.0.i386.rpm or DoubleTake-7.0.0.1124.0.x86 64.rpm.
 - 32-bit architecture—DoubleTake-version_numbers.i386.rpm
 - 64-bit architecture—DoubleTake-version_numbers.x86_64.rpm
- 2. If you are upgrading an existing Double-Take server, you should complete the following steps before upgrading.
 - Shutdown the protected application(s)
 - Stop the Double-Take daemon
 - Detach DTLOOP, if it is being used
- 3. Once you have determined the appropriate installation package to use, you can install the software from the UI or from the command line.
 - **UI installation**—Double-click the .rpm file from the UI and confirm the installation or upgrade.
 - Command line installation—Use the following steps to install from a command line.
 - a. Make sure you are a root/uid 0 user.
 - b. Go to a shell prompt by launching a terminal application from your UI or logging in via the Linux virtual console.
 - c. If you are installing from a CD, mount the file ISO 9660 or UDF file system.
 - d. To run the installation use rpm -i with the installation file name to install the software or rpm -U with the installation file name to upgrade the software. For example, if you were installing on a 32-bit operating system, you would use the command rpm -i DoubleTake-7.0.0.1124.0.i386.rpm.

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

Use DTSetup to configure your servers for Double-Take and restart the daemon if you upgraded. S DTSetup on page 30.	ee

Installing or upgrading Double-Take Windows client

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

- 1. Close any open applications.
- 2. Start the installation program using the appropriate instructions, depending on your media source.
 - Physical media—If auto-run is enabled, the installation program will start automatically.
 To manually start the program, select Start, Run and specify <drive>:\autorun.exe.
 - Web download—Launch the .exe file that you downloaded from the web.
- When the installation program begins, the Autorun appears allowing you to install software and view documentation and online resources. Select the Install Double-Take for Linux Management Client link.
- 4. When the Welcome screen is displayed. Click **Next** to continue.
- 5. Review and accept the Vision Solutions license agreement to continue with the installation program. Click **Next** to continue.
- Select the folder where you would like to install the Double-Take clients and click **Next** to continue.
- 7. When you are ready to begin copying the files, click **Install**.
- 8. 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. Make sure you are a root/uid 0 user.
- 2. 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.



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

Chapter 3 DTSetup

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

- Running DTSetup on page 31—This topic includes instructions for launching DTSetup.
- Setup tasks on page 32—The setup tasks allow you to configure activation codes, security groups, legacy replication configuration, server configuration, and driver performance settings.
- Starting and stopping the daemon on page 40—Built-in scripts allow you to quickly and easily start and stop the Double-Take daemon.
- Starting DTCL on page 41—You can launch the Double-Take interactive command prompt which allows you to enter DTCL commands one at a time.
- Viewing documentation and troubleshooting tools on page 42—DTSetup provides easy access to Double-Take log files, a diagnostic collection tool, and several legal documents.
- *DTSetup menus* on page 43—This topic includes a list overview of the DTSetup menu system. Reference the links in the list for complete details on completing tasks in DTSetup.

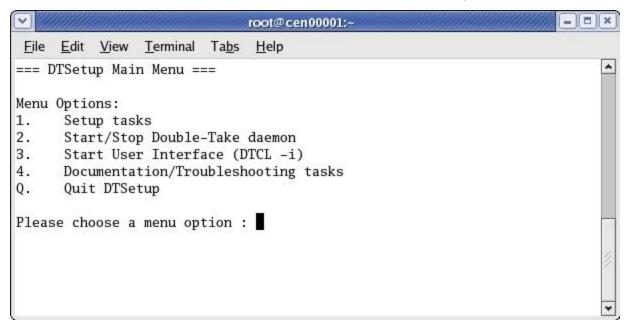
Running DTSetup

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



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

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



Setup tasks

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

- Activating your server on page 32—Activation codes license your Double-Take servers.
- Modifying security groups on page 33—Security groups provide access to Double-Take.
- Replication configuration on page 33—You must configure each source server for replication. You have the choose of selecting block device replication or legacy replication.
- Configuring server settings on page 38—If desired, you can modify server settings through the Double-Take configuration file.
- Configuring driver performance settings on page 39—If desired, you can specify Double-Take driver performance settings.

Activating your server

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



Server activation can also be completed through the Replication Console. See *Licensing a server* on page 190.

- 1. Start DTSetup. See Running DTSetup on page 31.
- 2. Select **Setup tasks**.
- 3. Select Set Activation Code Menu.
- 4. Select Set Activation Code in /etc/DT/DT.conf.
- 5. Enter your activation code and press Enter. The activation code will automatically be inserted into the configuration file. You are prompted to start the Double-Take service after the first installation, and you must restart the service each time the activation code is modified, such as after an upgrade.
- 6. Press Enter to return to the menu.
- 7. Press Q as many times as needed to return back to the main menu or to exit DTSetup.

Modifying security groups

During the installation, the user root is automatically added to the Double-Take administrators security group. If you want to add other users or remove root, you will need to modify the security group configuration for each source and target server. See *Security* on page 209 for more details on the security groups and the privileges granted to each group.

- 1. Start DTSetup. See Running DTSetup on page 31.
- 2. Select Setup tasks.
- 3. Select Add/Remove users to Double-Take groups.
- 4. Select the necessary menu options to add or remove groups to the administrator or monitors group as needed, and specify the user name when prompted.
- 5. When you have completed your security group modifications, press Q as many times as needed to return back to the main menu or to exit DTSetup.

Replication configuration

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

- Configuring file system replication on page 34—This option has been replaced by a new filter
 driver. You may be instructed to use this legacy driver by technical support. It uses a Double-Take
 file system mount point to capture data changes at the byte level. It provides a more detailed level
 of data selection, allowing you to include and exclude files and directories.
- Configuring block device replication on page 36—Block device replication uses a loop device to capture data changes at the block level. Data selection is at the volume level. Generally, this replication configuration is ideal for non-file system type data, like raw databases.

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

Configuring file system replication

This option has been replaced by a new filter driver. You may be instructed to use this legacy driver by technical support.

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. Data on the file system can only be replicated if it is under DTFS mount points.

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



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

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

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

- 1. Start DTSetup. See Running DTSetup on page 31.
- 2. Select Setup tasks.
- 3. Select Configure File System or Block Device Replication.
- 4. Select (DEPRECATED) Configure File System Replication setup file.
- 5. To add a DTFS mount, select Add an entry to /etc/DT/dtfs mounts.
- 6. Specify the path of the highest subdirectory for which replication should be captured. This can be the path to a mount point or a subdirectory of a volume.



Do not select the root as a replication mount point.

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

7. You will be prompted whether you want to replicate access times. The access-time file property changes during a mirror and causes replication operations to be generated during a mirror and every time a file is read from or written to. This may have a negative effect on performance. If you do not replicate access times on both the source and target when the volume is initially mounted, you may increase the performance, especially during mirroring. (Access times from files closes are always replicated regardless of this configuration.) Specify yes or no to replicate access times.

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



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

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



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

Configuring block device replication

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

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



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

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

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

- 1. Start DTSetup. See Running DTSetup on page 31.
- 2. Select **Setup tasks**.
- 3. Select Configure File System or Block Device Replication.
- 4. Select Configure Block Device Replication setup file.
- 5. If you want to see a list of block devices to which DTLOOP can be attached, select **List block devices on this system**. Swap and LVM physical partitions will not be included in the list.
- 6. Press Enter to continue.
- 7. To add a DTLOOP device, select Add an entry to /etc/DT/dtloop_devices.
- 8. Enter the path to the block device that is to be replicated and press Enter.
- 9. Enter the path to the loop device to use (/dev/loop#), if the same one should always be attached. If you use more than one loop device, you should assign a specific number to the loop device so it will persist beyond reboots. DTLOOP can also use the first one available, but that may mean it attaches to a different one on subsequent reboots/restarts, which may not be desirable. Press Enter to continue.
- 10. You will be asked if you want to attach at an offset into the block device and if you want to use an encrypted loop device. In general, these options can be left blank. See the losetup man page for more information on using encryption.
- 11. To remove any DTLOOP devices, select **Remove /Detach entries in /etc/DT/dtloop_devices** and specify the path that you want to remove.

- 12. To immediately attach all of the entries in /etc/DT/dtloop_devices, select **Attach all entries in** /etc/DT/dtloop_devices to a loop device. If you do not select this option, you must stop and restart the daemon for the changes to take effect.
- 13. To immediately detach all of the entries in /etc/DT/dtloop_devices, select **Detach loop devices** from all entries in /etc/DT/dtloop_devices. This allows you to make changes to DTLOOP without unloading the daemon.
- 14. When you have completed your block device replication configuration, press Q as many times as needed to return back to the main menu or to exit DTSetup.



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

Configuring server settings

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

- 1. Start DTSetup. See Running DTSetup on page 31.
- 2. Select Setup tasks.
- 3. Select Edit Double-Take config file.
- 4. The server settings are listed in alphabetical order. Make modifications as necessary, using the control keys specified at the bottom of the page. For a complete list of each server setting, valid values, default values, and optional notes, see *Server Settings* in the *Scripting Guide*.
- 5. Press control-X to exit the configuration file.
- 6. Enter Yes or No to save any changes.
- 7. Press Q as many times as needed to return back to the main menu or to exit DTSetup.

Configuring driver performance settings

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



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

- 1. Start DTSetup. See Running DTSetup on page 31.
- 2. Select Setup tasks.
- 3. Select Configure Double-Take driver performance.
- 4. The current driver settings are displayed.
- 5. Select a driver setting to modify the option.
 - Toggle Adaptive Throttling—You can toggle between enabling (true) and disabling
 (false) Adaptive Throttling. This occurs when kernel memory usage exceeds the
 Throttling Start Level percentage. When throttling is enabled, operations are delayed by,
 at most, the amount of time set in Maximum Throttling Delay, thus reducing kernel
 memory usage. Throttling stops when the kernel memory usage drops below the
 Throttling Stop Level percentage.
 - Toggle Forced Adaptive Throttling—You can toggle between enabling (true) and
 disabling (false) Forced Adaptive Throttling. This causes all operations to be delayed by,
 at most, the amount of time in set in Maximum Throttling Delay, regardless of the kernel
 memory being used. Adaptive Throttling must be enabled (true) in order for Forced
 Adaptive Throttling to work.
 - **Set Maximum Throttling Delay**—This option is the maximum time delay, in milliseconds, used by the driver during a system delay.
 - **Set Throttling Delay Interval**—This option is the interval, in milliseconds, to check memory usage during a throttling delay. If a delay is no longer needed, the remainder of the delay is skipped.
 - **Set Throttling Start Level**—Throttling starts when disk writes reach the specified percentage. This prevents the driver from stopping replication because memory has been exhausted.
 - **Set Throttling Stop Level**—Throttling stops when disk writes reach the specified percentage.
 - **Set Memory Usage Limit**—This option is the amount of kernel memory, in bytes, used for queuing replication operations. When this limit is exceeded, the driver will send an error to the daemon forcing a remirror of all active connections.
 - **Set Maximum Write Buffer Size**—This option is the maximum amount of system memory, in bytes, allowed for a single write operation. Operations exceeding this amount are split into separate operations in the queue.
- 6. After you have completed your driver performance modifications, press Q as many times as needed to return back to the main menu or to exit DTSetup.

Starting and stopping the daemon

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

- 1. Start DTSetup. See Running DTSetup on page 31.
- 2. Select Start/Stop Double-Take daemon.
- 3. Select the necessary menu option to start or stop the daemon and handle the driver configuration.
 - Start Double-Take and process driver config—This option starts the Double-Take daemon and loads the Double-Take drivers.
 - Stop Double-Take but preserve driver config—This option stops the Double-Take daemon but does not unload the Double-Take drivers.
 - Restart service but preserve driver config—This option does a full stop and start of the Double-Take daemon but does not unload the Double-Take drivers.
 - Restart service and reset driver config—This option does a full stop and start, completely unloading the Double-Take daemon and Double-Take drivers and then reloading them.
 - Stop the running service and teardown driver config—This option stops the Double-Take daemon and the Double-Take drivers are unloaded.
 - Go to Replication Configuration menu—This option takes you to Setup Tasks, Configure File System or Block Device Replication, (DEPRECATED) Manual Replication Configuration menu. When you press Q to exit from that menu, you will return this menu.
- 4. When you have completed your starting and stopping tasks, press Q as many times as needed to return back to the main menu or to exit DTSetup.

Starting DTCL

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

- 1. Start DTSetup. See Running DTSetup on page 31.
- 2. Select Start User Interface (DTCL -i).
- 3. Enter your DTCL commands one at a time at the **Command** prompt. For a complete list of DTCL commands, their syntax, and instructions for completing tasks using DTCL, see the *Scripting Guide*.
- 4. To exit the DTCL Command prompt, type exit.
- 5. When you have completed your DTCL tasks, press Q as many times as needed to return back to the main menu or to exit DTSetup.

Viewing documentation and troubleshooting tools

- 1. Start DTSetup. See Running DTSetup on page 31.
- 2. Select Documentation/Troubleshooting tasks.
- 3. Select **View log files** to view the following log files. Double-Take logs alerts, which are processing notifications, warnings, and error messages. The logs are written to disk.
 - **View** *I**.**dtl in less**—This option uses the less file viewer program to view all of the Double-Take logs, starting from the most recent.
 - **Follow the output of latest**—This option uses tail -f to watch the output of the Double-Take logs in real-time.
 - View /var/log/messages in less—This option uses the less file viewer program to view the system log messages.
 - Follow the output of /var/log/messages—This option uses tail -f to watch the output of the system log messages in real-time.
- 4. Select one of the Collect and package diagnostic info selections to run the DTInfo script which collects configuration data. This can be useful when reporting problems to technical support. Depending on the diagnostic option you select, the amount of data to be collected varies between basic, detailed and full diagnostic information. You must have root (or uid 0 equivalent) to execute the diagnostics or to copy or read the resulting file.
- 5. Select **View user documentation** to view several legal documents. DTSetup attempts to determine your viewers, although you can specify your viewer.
 - View End User License Agreement TXT—This option views the End User License Agreement legal document.
 - View driver module license TXT—This option views the open source legal document.
 - Change a document viewer—This option allows you to specify a document viewer.
- 6. When you have completed your documentation and troubleshooting tasks, press Q as many times as needed to return back to the main menu or to exit DTSetup.

DTSetup menus

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

- 1. **Setup tasks**—Activation codes, security groups, replication configuration, server configuration, and driver performance settings. See *Setup tasks* on page 32.
 - 1. **Set Activation Code Menu**—See *Activating your server* on page 32.
 - 2. Add/Remove users to Double-Take groups—See Modifying security groups on page 33
 - 3. **Configure File System or Block Device Replication**—There are three types of replication configurations for your source server. See *Replication configuration* on page 33.
 - 1. **(DEPRECATED) Configure File System Replication setup file**—See *Configuring file system replication* on page 34.
 - 2. **Configure Block Device Replication setup file**—See *Configuring block device replication* on page 36.
 - 3. (DEPRECATED) Manual Replication Configuration menu—See Configuring file system replication on page 34 and Configuring block device replication on page 36
 - 4. **(DEPRECATED) Full Server Replication Configuration menu**—This option is no longer being used.
 - 4. Edit Double-Take config file—See Configuring server settings on page 38.
 - 5. **Configure Double-Take driver performance**—See *Configuring driver performance settings* on page 39.
- 2. Start/Stop Double-Take daemon—See Starting and stopping the daemon on page 40.
- 3. Start User Interface (DTCL -i)—See Starting DTCL on page 41.
- 4. **Documentation/Troubleshooting tasks**—See *Viewing documentation and troubleshooting tools* on page 42.

Chapter 4 Clients

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



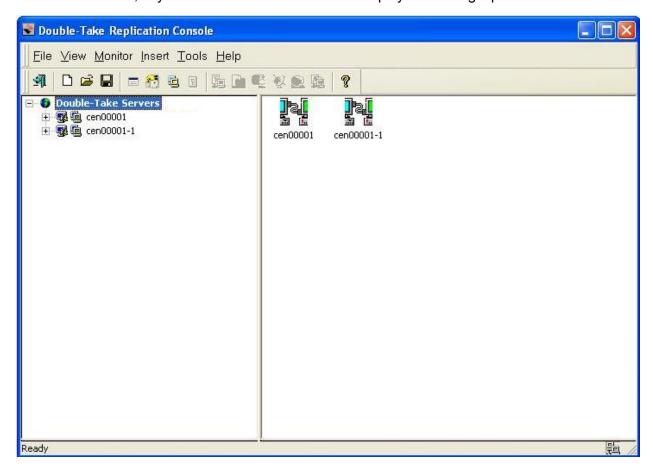
Double-Take also has a scripting language which can be used in an interactive client or in scripts. For more information, see the *Scripting Guide*.

- Replication Console on page 45
- Failover Control Center on page 48

Replication Console

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

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



Using Replication Console workspaces

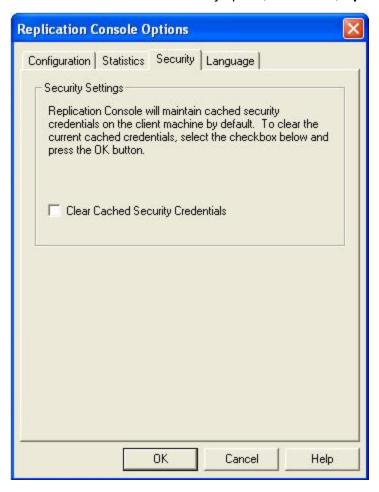
The Replication Console workspace contains the display of the panes of the Replication Console and any servers that may have been inserted. Multiple workspaces can be used to help organize your environment or to view settings from another machine.

- Saving a workspace—As you size, add, or remove windows in the Replication Console, you can save the workspace to use later or use on another Double-Take client machine. Select **File** and one of the following options.
- Save Workspace—Save the current workspace. If you have not previously saved this
 workspace, you must specify a name for this workspace.
- Save Workspace As—Prompt for a new name when saving the current workspace.
- Opening a workspace—From the Replication Console, you can open a new workspace or open a previously saved workspace. Select File and one of the following options.
- **New Workspace**—Open an untitled workspace with the default Double-Take window settings.
- Open Workspace—Open a previously saved workspace.

Clearing stored security credentials

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

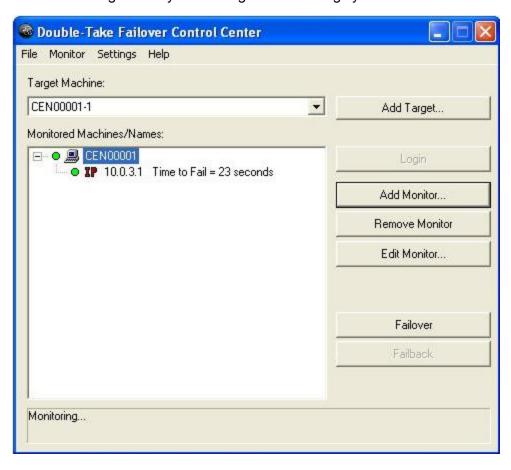
1. To access the credentials security option, select **File**, **Options** and select the **Security** tab.



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

Failover Control Center

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



Setting the frequency of Failover Control Center console refreshes

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

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

Chapter 5 Data protection

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

You have the following data protection options.

- Automated process—If you would like to use an automated process that walks you through both the replication and connection tasks, you only need to complete the Connection Wizard steps. See Establishing a data connection using the automated Connection Wizard on page 51.
- Manual process—If you want to go through the tasks manually, begin by creating a replication set and then continue with establishing a connection. See Creating a replication set on page 53 and Establishing a connection manually using the Connection Manager on page 56.
- NAT or firewall—If your environment has a NAT or firewall configuration, you will need to begin
 with creating a replication set and then follow the instructions for establishing a NAT connection.
 See Creating a replication set on page 53 and Establishing a connection across a NAT or firewall
 on page 60.
- Simulating a connection—If you want to simulate a connection for planning purposes, begin by
 creating a replication set and then continue with establishing a simulated connection. See
 Creating a replication set on page 53 and Simulating a connection on page 62.

Establishing a data connection using the automated Connection Wizard

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

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



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

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



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

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



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

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



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

- 6. Click **Next** to continue.
- 7. Select **Protect data** and click **Next** to continue.
- 8. Choose to create a new replication set or use a replication set that already exists.
 - Create a new replication set with this name—If you choose to create a new replication, specify a replication set name.
 - **Use this replication set**—If you choose to use an existing replication set, specify the name of that replication set by selecting it from the pull-down menu.

- 9. Click **Next** to continue.
- 10. If you are creating a new replication set, a tree display appears identifying the volumes and directories available on your selected source server. Mark the check box of the volumes and/or directories you want to protect and click **Next** to continue.
- 11. Select the location on the target where the data will be stored.
 - Send all data to a single path on the target—This option sends all selected volumes and directories to the same location on the target. The default location is /source_ name/replication set name/volume name.
 - Send all data to the same path on the target—This option sends all selected volumes and directories to the same directories on the target.
 - **Custom**—To select a custom path, click once in the Target Path field and modify the drive and directory to the desired location.
- 12. Click **Next** to continue.
- 13. Review your selections on the summary screen. If your Connection Wizard settings are correct, establish your connection by completing one of the two options below.
 - If you do not want to set advanced options, click **Finish**. The Connection Wizard will close, the connection will be established, and mirroring and replication will begin.
 - If you want to set advanced options, click Advanced Options. The Connection Wizard will close and the Double-Take Connection Manager will open. The Servers tab will be completed.

Creating a replication set

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

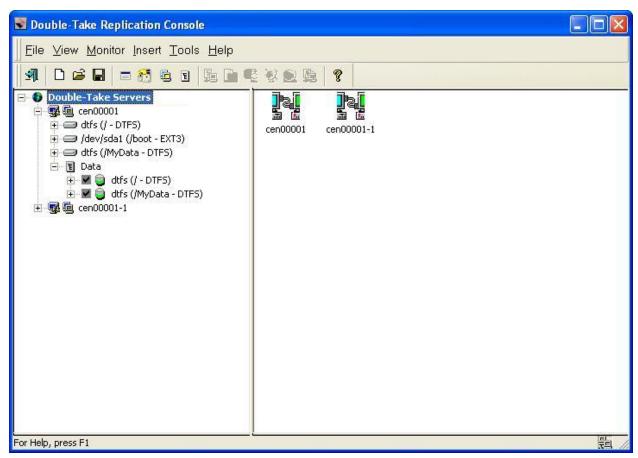
- Highlight a source in the left pane of the Replication Console and select Insert, Replication Set from the menu bar. You can also right-click on the source name and select New, Replication Set.
- 2. A replication set icon appears in the left pane under the source. By default, it is named New Replication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



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

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

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





Be sure and verify what files can be included by reviewing the *Replication capabilities* on page 132.

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

- 5. After selecting the data for this replication set, right-click the new replication set icon and select **Save**. A saved replication set icon will change from red to black.
- 6. If you need to select a block device for replication, right-click the replication set and select **Add Device**.
- 7. The block devices configured for Double-Take Availability replication are shown by default. Highlight the device to include in the replication set and click **OK**.



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



any of these devices, but you cannot use them for Double-Take Availability replication until they are configured for Double-Take Availability replication. The status **no dtloop** indicates the device is not configured for Double-Take Availability replication.

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

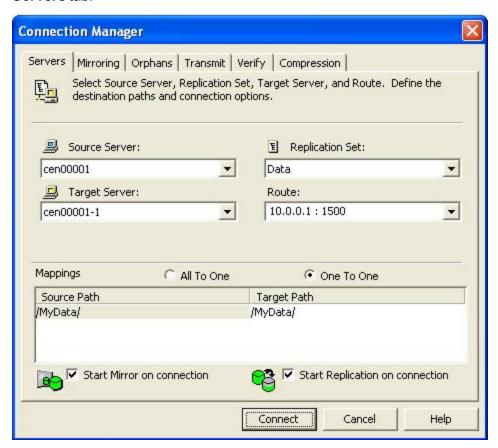
The partition size displayed may not match the output of the Linux df command. This is because df shows the size of the mounted file system not the underlying partition which may be larger. Additionally, Double-Take Availability uses powers of 1024 when computing GB, MB, and so on. The df command typically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

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

Establishing a connection manually using the Connection Manager

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

- 1. Open the Connection Manager to establish the connection.
 - Highlight the replication set and select Tools, Connection Manager.
 - Right-click on the replication set and select Connection Manager.
 - Drag and drop the replication set onto a target. The target icon could be in the left or right pane of the Replication Console.
- 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.



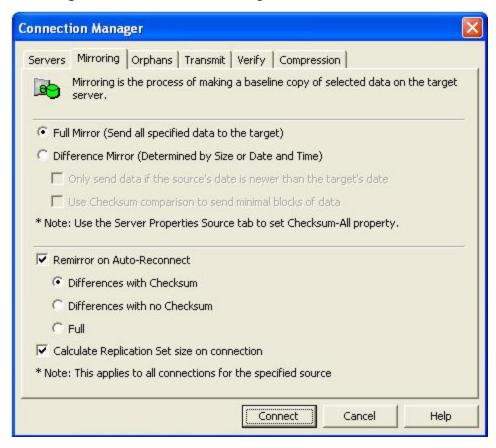
- **Source Server**—Specify the source server that contains the replication set that is going to be transmitted to the Double-Take target.
- Replication Set—At least one replication set must exist on the source before establishing
 a connection. Specify the replication set that will be connected to the 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 DoubleTake 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.



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

• Start Replication on Connection—Replication can be initiated immediately when the connection is established. If replication is not configured to start automatically, you must start it manually after the connection is established. If you disable this option, you will need to perform a mirror prior to beginning replication to guarantee integrity.

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



- Full Mirror—All files in the replication set will be sent from the source to the target.
- **Difference Mirror**—Only those files that are different based size or date and time (depending on files or block devices) will be sent from the source to the target.
 - Only send data if the source's date is newer than the target's date—Only those files that are newer on the source are sent to the target.



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

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



See Stopping, starting, pausing, or resuming mirroring on page 125 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.



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

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

Establishing a connection across a NAT or firewall

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

In this environment, you must have static mapping where a single, internal IP address is always mapped in a one-to-one correlation to a single, external IP address. Double-Take 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.

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

Replication Console Status Transmit Port

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

Replication Console Heartbeat Advertisement

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

Failover Control Center Heartbeat Listen Port

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

Double-Take Server Heartbeat Transmit Port

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

Double-Take Server Status Listen Port

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

Double-Take Server Statistics Logging Port

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

- 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.

- 3. If your network is configured to propagate UDP broadcasts, your servers will be populated in the Replication Console from across the router. If not, you have to manually insert the servers, by selecting **Insert**, **Server**. Type the IP address of the router the server is connected to and the port number the server is using for heartbeats.
- 4. Once your server is inserted in the Replication Console, you can use the Connection Wizard or the Connection Manager to establish your connection. See *Establishing a data connection using the automated Connection Wizard* on page 51 or *Establishing a connection manually using the Connection Manager* on page 56.

Simulating a connection

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.

- 1. Before and after simulating your connection, you should gather network and system information specific to Double-Take operations. Use DTSetup to run DTInfo to automatically collect this data.
- 2. Select the DTSetup option for troubleshooting, then select the option for basic diagnostics.
- 3. When you run the diagnostics, it may take several minutes for it to finish processing. When it is complete, a .tar.gz file will be created in /var/run/etc/DT/. The file name will have DTInfo with the date and time. You must have root (or uid 0 equivalent) to execute the diagnostics or to copy or read the resulting file.
- 4. Opening the Connection Manager to establish the connection.
 - Highlight the replication set and select Tools, Connection Manager.
 - Right-click on the replication set and select Connection Manager.
- 5. The Connection Manager opens to the Servers tab. Depending on how you opened the Connection Manager, some entries on the Servers tab will be completed already. For example, if you accessed the Connection Manager by right-clicking on a replication set, the name of the replication set will be displayed in the Connection Manager. Verify or complete the fields on the Servers tab.
 - Source Server—Specify the source server that contains the replication set that is going to be simulated to the TDU.
 - Replication Set—At least one replication set must exist on the source before establishing
 a connection. Specify the replication set that will be connected to the TDU.
 - Target Server—Select the Diagnostics target.
 - Route—After selecting the Diagnostics target, the Route will automatically be populated with Throughput Diagnostics Utility (TDU).
 - **Mappings**—Mappings are not required when simulating a connection because no data is actually transmitted to the target.
 - Start Mirror on Connection—Make sure this option is selected so that your simulation will be realistic.
 - Start Replication on Connection—Make sure this option is selected so that your simulation will be realistic.
- Click Connect to establish the connection. The simulation data will be logged to the Double-Take statistics file.
- 7. Repeat steps 1-3 to run the diagnostics utility after the simulation is complete.

Chapter 6 Protection monitoring

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

- Monitoring a data workload on page 64—You can use the build-in monitoring in the Replication Console to watch statistics and see at-a-glance the health of your protection.
- Log files on page 70—The Double-Take log files provide notification, warning, and error processing messages.
- Monitoring the Linux system log on page 81—Double-Take generates Linux system messages.
- Statistics on page 95—Additional statistics are available outside of the Replication Console.
- SNMP on page 105—Both statistics and processing messages are available through SNMP.

Monitoring a data workload

When a source is highlighted in the left pane of the Replication Console, the connections and their statistics are displayed in the right pane. Additionally, colors and icons are used for the connections, and the Double-Take servers, to help you monitor your connections.

- Connection statistics on page 64
- Connection and sever display on page 68

Connection statistics

- 1. You can change the statistics that are displayed by selecting **File**, **Options** and selecting the **Statistics** tab.
- 2. The statistics displayed in the Replication Console will be listed with check boxes to the left of each item. Mark the check box to the left of each statistic that you want to appear, and clear the check box to the left of each statistic that you do not want to appear.
- 3. The statistics appear on the Replication Console in the order they appear on the **Statistics** tab. If you want to reorder the statistics, highlight the statistic to be moved and select the up or down arrow button, to the right of the vertical scroll bar, to move the selection up or down in the list. Repeat this process for each statistic that needs to be moved until you reach the desired order.
- 4. If you have made changes to the statistics list and have not yet saved them, you can go back to the previously used settings by clicking **Reset to Last**. This will revert the list back to the last saved settings.
- 5. To return the statistics list to the Double-Take default selection and order, click **Reset to Default**.
- 6. Click **OK** to apply and save any changes that have been made to the order or display of the Replication Console statistics.

Statistics marked with an asterisk (*) are not displayed, by default.

Replication Set

Replication set indicates the name of the connected replication set.

Connection ID

The connection ID is the incremental counter used to number each connection established. This number is reset to one each time the Double-Take service is restarted.

Target Name

The name of the target as it appears in the server tree in the left pane of the Replication Console. If the server's name is not in the server tree, the IP address will be displayed.

Target IP

The target IP is the IP address on the target machine where the mirroring and replication data is being transmitted.

Target Data State

- OK—The data on the target is in a good state.
- **Mirroring**—The target is in the middle of a mirror process. The data will not be in a good state until the mirror is complete.
- Mirror Required—The data on the target is not in a good state because a remirror is required. This may be caused by an incomplete or stopped mirror or an operation may have been dropped on the target.
- Restore required—The data on the source and target do not match because of a
 failover condition. Restore the data from the target back to the source. If you want to
 discard the changes on the target, you can remirror to resynchronize the source and
 target.
- Not Ready—The Linux drivers have not yet completed loading on the target.

Target Status

- OK—The target machine is active and online.
- Not Loaded—The target module is not loaded on the target. (For example, the
 activation code is invalid.)
- Paused—The target machine is paused by user intervention.
- Retrying—The target machine is retrying operations for the connection.

This field may not be updated until there is source/target activity.

Commit Mode *

The commit mode status indicates the connection status.

- Real-time—Data is being transmitted to the target machine in real-time.
- **Scheduled**—Data is waiting to be transmitted to the target machine until one or more transmit options have been met.

Transmit Mode

- Started—Data is being transferred to the target machine.
- Paused—If the transmission is real-time and the transmission has been paused, the Transmit Mode indicates Paused.
- Scheduled—If the transmission is scheduled, the Transmit Mode indicates Scheduled.
- **Stopped**—Data is not being transferred to the target machine.
- Error—There is a transmission error.

Mirror Status

- Mirroring—If the file size of the replication set has not been calculated and the data
 is being mirrored to the target machine, the Mirror Status will indicate Mirroring.
- Idle—Data is not being mirrored to the target machine.
- Paused—Mirroring has been paused.
- Percentage Complete—If the file size of the replication set has been calculated
 and the data is being mirrored to the target machine, the Mirror Status will display
 the percentage of the replication set that has been sent.

- Waiting—Mirroring is complete, but data is still being written to the target.
- Restoring—Data is being restored from the target to the source.
- Verifying—Data is being verified.
- Removing Orphans—Double-Take is checking for orphan files within the target path location (files that exist on the target but not on the source). These files will be removed.

Replication Status

- Replicating—Data is being replicated to the target machine.
- Ready—There is no data to replicate to the target machine.
- Stopped—Replication has stopped.
- Pending—If auto-remirror is enabled and you have experienced a source or target failure and recovery, the status will change to pending while the connections are reestablished and will update when the remirror begins. If auto-remirror is disabled and you have experienced a source or target failure and recovery, replication will be Pending until a remirror is performed. Without a remirror, data integrity cannot be guaranteed.
- Out of Memory—Kernel memory has been exhausted.

Queued (Ops) *

The queued (ops) statistic indicates the total number of mirror and replication operations that are in the source queue.

Sent (Bytes)

The sent (bytes) statistic indicates the total number of mirror and replication bytes that have been transmitted to the target.

Sent Compressed (Bytes)

The sent compressed (bytes) statistic indicates the total number of compressed mirror and replication bytes that have been transmitted to the target. If compression is disabled, this statistic will be the same as sent (bytes).

Intermediate Queue (Bytes) *

The intermediate queue (bytes) indicates the total amount of memory being used by the operations buffer queue.

Disk Queue (Bytes)

The disk queue (bytes) indicates the amount of disk being used to queue data on the source.

Queued Replication (Bytes)

The queued replication (bytes) statistic is the total number of replication bytes that are remaining to be transmitted from the source.

Sent Replication (Bytes)

The sent replication (bytes) statistic is the total number of replication bytes that have been transmitted to the target.

Sent Compressed Replication (Bytes) *

The sent compressed replication (bytes) statistic is the total number of compressed replication bytes that have been transmitted to the target. If compression is disabled, this statistic will be the same as sent replication (bytes).

Queued Mirror (Ops) *

The queue mirror (ops) statistic is the total number of mirror operations in the queue.

Sent Mirror (Bytes)

The sent mirror (bytes) statistic is the total number of mirror bytes that have been transmitted to the target.

Sent Compressed Mirror (Bytes) *

The sent compressed mirror (bytes) statistic is the total number of compressed mirror bytes that have been transmitted to the target. If compression is disabled, this statistic will be the same as sent mirror (bytes).

Skipped Mirror (Bytes)

The skipped mirror (bytes) statistic is the total number of bytes that have been skipped when performing a difference or checksum mirror. These bytes are skipped because the data is not different on the source and target machines.

Remaining Mirror (Bytes)

The remaining mirror (bytes) statistic is the total number of mirror bytes that are remaining to be sent to the target.

Queued Replication (Ops) *

The queued replication (ops) statistic is the total number of replication operations in the queue.

Last File Touched

The last file touched identifies the last file that Double-Take transmitted to the target. If you are using long file names (more than several thousand characters long) you may want to disable the display of this statistic to improve Replication Console response times.

Connected Since

Connected since is the date and time indicating when the current connection was made. This field is blank, indicating that a TCP/IP socket is not present, when the connection is waiting on transmit options or if the transmission has been stopped. This field will maintain the date and time, indicating that a TCP/IP socket is present, when transmission has been paused.

Connection and sever display

You can configure when the icons and colors change to accommodate your network environment. For example, a slow or busy network may need longer delays before updating the icons or colors.

- Select File, Options. On the Configuration tab, you will see Site Monitor and Connection Monitor. The Site Monitor fields control the icons on the left pane of the Replication Console and the icons on the right pane when a group is highlighted in the left pane. The Connection Monitor field controls the display when a server is highlighted in the left pane. These two separate monitoring capabilities allow for flexible monitoring.
- 2. Under **Site Monitor**, specify **Check Status Interval** to identify the number of seconds between requests sent from the Replication Console to the servers in order to update the display. Valid values are between 0 and 3600. The default setting is 30 seconds.
- 3. Under **Site Monitor**, specify **Missed Status Responses** to identify the number of responses from a server that can be missed before the Replication Console considers communications lost and updates the icons. Valid values are between 1 and 100. The default setting is 2.
- 4. Under Connection Monitor, specify Missed Status Responses to identify the number of responses from a server that can be missed before the Replication Console considers communications lost and updates the icons and colors. Valid values are between 0 and 1000. The default setting is 5.
- 5. Click **OK** to save the settings.



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

The following icons are displayed in the left pane.

- —An icon with yellow and blue servers indicates a server that is working properly.
- —A red X on a server icon indicates the Replication Console cannot communicate with that server or that is a problem with one of the server's connections. If the connection background is gray, it is a communication issue. If the connection also has a red X, it is a connection issue.
- —A red tree view (folder structure) on a server icon indicates a restore is required because of a failover.
- △ → A black X on a server icon indicates the server is not running Double-Take.

The following icons and colors are displayed in the right pane when a server is highlighted in the left pane.

—A green checkmark on a connection indicates the connection is working properly.

➤ —A red X on a connection indicates a connection error. For example, an error may be caused by broken transmission or pending replication. To determine the exact problem, locate the connection data item that appears in red.

White background—If the connection background is white, the Replication Console and the source are communicating.

Gray background—If the connection background is gray, the Replication Console and the source are no longer communicating. The connection data stops updating once communications have stopped. Once communications have been reestablished, the connection background will change back to white.

Log files

Various Double-Take components (Double-Take service, Replication Console, Failover Control Center, and the Command Line Client) generate a log file to gather alerts, which are notification, warning, and error messages. The log files are written to disk.

Each log file consists of a base name, a series number, and an extension.

- Base Name—The base name is determined by the application or process that is running.
 - Double-Take—dtlog
 - Replication Console—mc
 - Failover Control Center—fcc
 - Command Line Client—dtcl
- Series Number—The series number ranges from 1 to 999. For example, Double-Take 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.
 - Double-Take—dtlog1.dtl, dtlog2.dtl
 - Replication Console—mc1.dtl, mc2.dtl
 - Failover Control Center—fcc1.dtl, fcc2.dtl
 - Command Line Client—dtcl1.dtl, dtcl2.dtl

The following topics are available for the log file.

- Viewing the log files through a text editor on page 71
- Viewing the Double-Take log file through the Replication Console on page 72
- Configuring the properties of the Double-Take log file on page 74
- Double-Take log messages on page 75

Viewing the log files through a text editor

The log files can be viewed, from the location where Double-Take is installed, with a standard text editor. The following list describes the information found in each column of the log file.

- 1. Date the message was generated
- 2. Time the message was generated
- Process ID
- 4. Thread ID
- 5. Sequence number is an incremental counter that assigns a unique number to each message
- 6. The type or level of message displayed 1 for warning or error message and 2 for informational message
- 7. Message ID
- 8. Message text

Sample Double-Take log file

```
01/15/2010 14:14:18.3900 95 98 2 2 69 Kernel Started
01/15/2010 14:14:18.4200 95 98 3 2 10004 Valid Activation Key Detected:
01/15/2010 14:14:18.5350 98 170 4 2 52501 Target module loaded successfully
01/15/2010 14:14:18.6760 98 172 5 2 10004 Valid Activation Key Detected:
01/15/2010 14:14:18.9870 130 131 6 2 51501 Source module loaded successfully
01/15/2010 14:24:15.2070 130 132 7 2 72 Connection Request from ip://206.31.4.305
01/15/2010 14:24:16.3090 131 133 8 2 600002 Unified login provides ADMIN access
01/15/2010 14:24:40.9680 132 134 9 2 99 RepSet Modified: UserData
01/15/2010 14:25:22.4070 134 131 10 2 71 Originator Attempting ip://206.31.4.305 01/15/2010 14:25:22.5030
134 131 11 2 0 Transmission Create to ip://206.31.4.305.
01/15/2010 14:25:22.6060 135 133 12 2 500000 UserData is connected to ip://206.31.4.305 01/15/2010
14:25:23.5030 136 98 13 2 87 Start Replication on connection 1
```

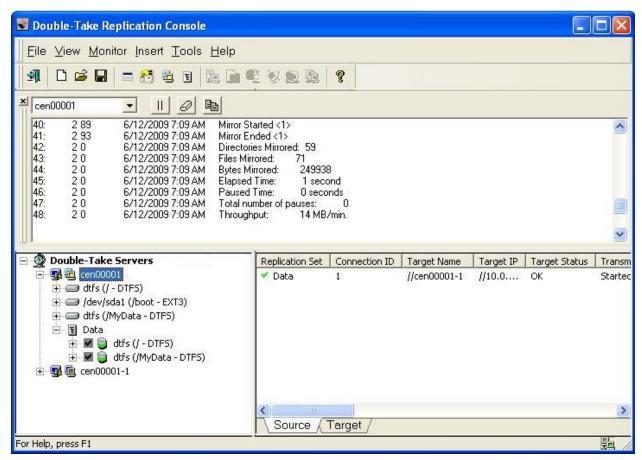
Sample Replication Console log file

```
00/00/0000 00:00:00.0000 Application starting
09/11/2010 12:45:53.8980 704 1032 1 2 0 Could not find XML file: C:\Program Files\Vision
Solutions\Double-Take for Linux\Administrator.xml, default groups will be added.
09/11/2010 12:45:53.9580 704 1032 2 2 0 Adding default group: Double-Take Servers
09/11/2010 12:45:53.9580 704 1032 3 2 0 Adding default group: Double-Take Servers\
Auto-Discovered Servers
09/11/2010 12:46:08.3390 704 1032 4 210004 Evaluation license expires in 95 day(s).
```

Viewing the Double-Take log file through the Replication Console

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

- 1. Open a new message window using any of the following methods.
 - Right-click on the server that you want to monitor in the left pane and select New, Message Window.
 - Select the Message Window icon from the toolbar.
 - Select Monitor, New Message Window and identify the Server that you want to monitor.
- 2. Repeat step 1 if you want to open multiple message windows.





The standard appearance of the message window is a white background. If your message window has a gray background, the window is inactive. The Replication Console may have lost communications with that server, for example, or you may no longer be logged into that server.

The message window is limited to the most recent 1000 lines. If any data is missing an entry in red will indicate the missing data. Regardless of the state of the message window, all data is maintained in the Double-Take log on the server.

- 3. To control the window after it is created, use one of the following methods to access the control methods listed in the following table.
 - Right-click on the message window and select the appropriate control.
 - Select the appropriate toolbar control.
 - Select Monitor, the name of the message window, and the appropriate control.

Close ×

Closes the message window



Clears the message window

Pause/Resume



Pauses and resumes the message window.

Pausing prevents new messages from being displayed in the message window so that you are not returned to the bottom of the message window every time a new message arrives. The messages that occur while the window is logged are still logged to the Double-Take 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.



Allows you to copy selected text

Options

This control is only available from the **Monitor** menu. Currently, there are no filter options available so this option only allows you to select a different server. In the future, this control will allow you to filter which messages to display.

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

Configuring the properties of the Double-Take log file

- 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 Replication Console and selecting **Properties**.
- 2. Select the **Logging** tab.
- 3. At the top of the window, **Folder** indicates the directory where the log files are located. The default is the directory where the Double-Take program files are installed.
- 4. Modify any of the options under **Messages and Alerts**, if necessary.
 - **Maximum Length**—Specify the maximum length of the log file. The default size is 1048576 bytes and is limited by the available hard drive space.
 - **Maximum Files**—Specify the maximum number of log files that are maintained. The default is 5 and the maximum is 999.



If you change the **Maximum Length** or **Maximum Files**, you must restart the Double-Take daemon for the change to take effect.

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

Double-Take log messages

The following list 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.



In this information, con_id refers to the unique connection ID assigned to each connection between a source replication set and a target.

There are several log messages with the ID of 0. See the description in the Message column in the log file.

7 Synchronous ioctl returned STATUS_PENDING

7 Failed to reset Replication Flags. Replication may not be performed correctly.

- Communication with the Double-Take 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 to undo the changes to the replication set, stop and restart Double-Take, or reboot the server.

69 Double-Take kernel started on server_name

The Double-Take service was started on the Double-Take server specified.

70 Double-Take kernel stopped

The Double-Take service was stopped on a Double-Take server.

71 Originator attempting ip://xxx.xxx.xxx

A source is requesting to connect a replication set to a target machine.

72 Connection request from ip://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

The transmission of data from the source machine to the target machine has resumed.

76 Connection failed to ip://xxx.xxx.xxx.xxx

An attempt to establish a network connection between a source machine and target machine has failed. Check your network connections and verify that the target machine is still online.

77 Connection lost with IP address address

The network connection previously established between a source machine and target machine has been lost. Check your network connections and troubleshoot to see why the connection was lost.

78 Auto-disconnect threshold has been reached.

The Double-Take queue has exceeded its limit, and the auto-disconnect process will disconnect the source and target connection. The auto-reconnect process will automatically reestablish the connection if the auto-reconnect feature is enabled. If the auto-reconnect feature is not enabled, you must first verify that the connection between the source and target has been broken, and then manually reestablish the connection in the Replication Console.

79 Memory freed to bring Double-Take memory usage below the limit

Data in the source queue has been sent to the target machine, bringing the pagefile below its limit.

80 Trying to auto-retransmit to ip://xxx.xxx.xxx.xxx

Double-Take is attempting to automatically reconnect previously established source and target connections after a server reboot or auto-disconnect. This is also referred to as the auto-reconnect process.

81 Schedule transmit start to target

A scheduled transmission of data from a source machine to a target machine has started. See the description in the Message column in the log file.

82 Schedule transmit end to target

A scheduled transmission of data from a source machine to a target machine has ended. See the description in the Message column in the log file.

85 repset has been auto-disconnected

Double-Take 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.

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 with connection con_id Disconnected replication set name: repset_name.

Auto-disconnect has occurred for the specified connection. This is due to the source queue filling up because of a network or target failure or bottleneck.

10001 Activation key is not valid.

An invalid activation code was identified when the Double-Take service was started.

10002 Evaluation period has expired.

The evaluation license has expired.

10003 Activation code violation with machine machine_name

Duplicate single-server activation codes are being used on the servers, and Double-Take is disabled.

10004 Valid activation key detected

A valid activation code was identified when the Double-Take service was started.

51001 Source module failed to load

The Double-Take 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.

52000 The target has been resumed due to manual intervention

The target has been paused or resumed through user intervention.

52000 Unfinished Op error

This error message contains various Microsoft API codes. The text Code -<x> Internal <y> appears at the end of this message. The code value indicates why the operation

failed, and the internal value indicates the type of operation that failed. These are the most common code values that appear in this error message.

- (5) Permission denied: The account running the Double-Take service does not have permission to update the file specified.
- (32) Sharing violation: Another application is using a particular file that Double-Take 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 service 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 Replication Console.

99999 Application is terminating normally

The Double-Take service is shutting down normally.

503010 Asyncloctl for status thread 178 terminated, terminating the status thread

A Double-Take process monitors the state of the Double-Take driver. When the Double-Take service is shut down, the driver is shut down, and this process is terminated. (No user action required.)

600002 Unified login provides ADMIN access

600002 User has level access (x)

- Using the current login grants ADMIN access.
- The listed user has listed access level and access level ID.

700000 The source machine source_machine is not responding to a ping.

This occurs when all monitored IP addresses on the source machine stop responding to pings. Countdown to failover will begin at the first occurrence and will continue until the source machine responds or until failover occurs.

800000 Active Directory GetHostSpns function call failed

800000 Active Directory RemoveSpns function call failed

800000 Active Directory AddSpns function call failed

- Double-Take 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.

Monitoring the Linux system log

An event is a significant occurrence in the system or in an application that requires administrators to be notified. The operating system writes notifications for these events to the Linux system log. The location of the log file depends on the configuration of /etc/syslog.conf, however, by default, it is /var/log/messages. The following table identifies the events generated by Double-Take.

1 This evaluation period has expired. Mirroring and replication have been stopped. To obtain a license, please contact your vendor.

Error—Contact your vendor to purchase either a single or site license.

2 The evaluation period expires in %1 day(s).

Information—Contact your vendor before the evaluation period expires to purchase either a single or site license.

3 The evaluation period has been activated and expires in %1 day(s).

Information—Contact your vendor before the evaluation period expires to purchase either a single or site license.

4 Duplicate activation codes detected on machine %1 from machine %2.

Warning—If you have an evaluation license or a site license, no action is necessary. If you have a single license, you must purchase either another single license or a site license.

5 This product edition can only be run on Windows Server or Advanced Server running the Server Appliance Kit.

Error—Verify your activation code has been entered correctly and contact technical support.

3000 Logger service was successfully started.

Information—No action required.

3001 Logger service was successfully stopped.

Information—No action required.

4000 Kernel was successfully started.

Information—No action required.

4001 Target service was successfully started.

Information—No action required.

4002 Source service was successfully started.

Information—No action required.

4003 Source service was successfully stopped.

Information—No action required.

4004 Target service was successfully stopped.

Information—No action required.

4005 Kernel was successfully stopped.

Information—No action required.

4006 Service has aborted due to the following unrecoverable error: %1

Error—Restart the Double-Take service.

4007 Auto-disconnecting from %1 (%2) for Replication Set %3, ID: %4 due to %5

Warning—The connection is auto-disconnecting because the disk-based queue on the source has been filled, the service has encountered an unknown file ID, the target server has restarted, or an error has occurred during disk queuing on the source or target (for example, Double-Take cannot read from or write to the transaction log file).

4008 Auto-disconnect has succeeded for %1 (%2) for Replication Set %3, ID: %4

Information—No action required.

4009 Auto-reconnecting Replication Set %1 to %2 (%3)

Information—No action required.

4010 Auto-reconnect has succeeded connecting Replication Set %1 to %2 (%3)

Information—No action required.

4011 Auto-reconnect has failed connecting Replication Set %1 to %2 (%3)

Error—Manually reestablish the replication set to target connection.

4014 Service has started network transmission.

Information—No action required.

4015 Service has stopped network transmission.

Information—No action required.

4016 Service has established a connection to %1 (%2) for Replication Set %3, ID: %4

Information—No action required.

4017 Service has disconnected from %1 (%2) for Replication Set %3, ID: %4

Information—No action required.

4018 %1, however, mirroring and replication have been disabled as a restore is required due to a previous failover.

Warning—Perform a restoration.

- 4019 Service has started a mirror to %1 (%2) for Replication Set %3, ID: %4

 Information—No action required.
- 4020 Service has paused a mirror to %1 (%2) for Replication Set %3, ID: %4

 Information—No action required.
- 4021 Service has resumed a mirror to %1 (%2) for Replication Set %3, ID: %4

 Information—No action required.
- 4022 Service has stopped a mirror to %1 for Replication Set %2, ID: %3, %4
 Information—No action required.
- 4023 Service has completed a mirror to %1 %2 for Replication Set %3, ID: %4, %5 Information—No action required.
- 4024 Service has started Replication to %1 (%2) for Replication Set %3, ID: %4

 Information—No action required.
- 4025 Service has stopped Replication to %1 (%2) for Replication Set %3, ID: %4

 Information—No action required.
- 4026 The target has been paused due to user intervention.

Information—No action required.

4027 The target has been resumed due to user intervention.

Information—No action required.

4028 Registration of service class with Active Directory failed. Verify that the Active Directory server is up and the service has the proper permissions to update its entries.

Warning—Verify that the Active Directory server is running and that the Double-Take service has permission to update Active Directory.

4029 Registration of service instance with Active Directory failed. Verify that the Active Directory server is up and the service has the proper permissions to update its entries.

Warning—Verify that the Active Directory server is running and that the Double-Take service has permission to update Active Directory.

4030 RSResource.dll has an unknown error. The product functionality has been disabled.

Error—Reinstall the software, using the installation Repair option, to install a new copy of the RSResource.dll. Contact technical support if this error persists.

4031 RSResource.dll could not be opened. The product functionality has been disabled.

Error—Reinstall the software, using the installation Repair option, to install a new copy of the RSResource.dll. Contact technical support if this error persists.

4032 The RSResource.dll component version does not match the component version expected by the product. The product functionality has been disabled.

Error—Reinstall the software, using the installation Repair option, to install a new copy of the RSResource.dll. Contact technical support if this error persists.

4033 RSResource.dll build version is invalid. The product functionality has been disabled.

Error—Reinstall the software, using the installation Repair option, to install a new copy of the RSResource.dll. Contact technical support if this error persists.

4034 Error verifying the service name. The product functionality has been disabled.

Error—Reinstall the software, using the installation Repair option, to install a new copy of the RSResource.dll. Contact technical support if this error persists.

4035 Error verifying the product name. The product functionality has been disabled.

Error—Reinstall the software, using the installation Repair option, to install a new copy of the RSResource.dll. Contact technical support if this error persists.

4036 Error verifying the vendor name. The product functionality has been disabled.

Error—Reinstall the software, using the installation Repair option, to install a new copy of the RSResource.dll. Contact technical support if this error persists.

4037 Error verifying the vendor URL name. The product functionality has been disabled.

Error—Reinstall the software, using the installation Repair option, to install a new copy of the RSResource.dll. Contact technical support if this error persists.

4038 Error verifying the product code. The product functionality has been disabled.

Error—Reinstall the software, using the installation Repair option, to install a new copy of the RSResource.dll. Contact technical support if this error persists.

4039 Error while reading RSResource.dll. The product functionality has been disabled.

Error—Reinstall the software, using the installation Repair option, to install a new copy of the RSResource.dll. Contact technical support if this error persists.

4040 The product code is illegal for this computer hardware. The product functionality has been disabled.

Error—Reinstall the software, using the installation Repair option, to install a new copy of the RSResource.dll. Contact technical support if this error persists.

4041 The product code is illegal for this operating system version. The product functionality has been disabled.

Error—Reinstall the software, using the installation Repair option, to install a new copy of the RSResource.dll. Contact technical support if this error persists.

4042 The product code requires installing the Windows Server Appliance Kit. The product functionality has been disabled.

Error—Reinstall the software, using the installation Repair option, to install a new copy of the RSResource.dll. Contact technical support if this error persists.

4043 This product can only be run on a limited number of processors and this server exceeds the limit. The product functionality has been disabled.

Error—Reinstall the software, using the installation Repair option, to install a new copy of the RSResource.dll. Contact technical support if this error persists.

4044 An error was encountered and replication has been stopped. It is necessary to stop and restart the service to correct this error.

Error—Contact technical support if this error persists.

4045 %1 value must be between 1025 and 65535. Using default of %2.

Error—Verify that the Double-Take port value you are trying to use is within the valid range. If it is not, it will automatically be reset to the default value.

4046 This service failed to start because of a possible port conflict. Win32 error: %1

Error—Verify that the Double-Take ports are not conflicting with ports used by other applications.

4047 Could not load ZLIB DLL %1. Some levels of compression will not be available.

Error—The compression levels available depend on your operating system. You can reinstall the software, using the installation Repair option, to install a new copy of the DynaZip.dll, or contact technical support if this error persists.

- 4048 Service has started a delete orphans task to %1 (%2) for Replication Set %3, ID: %4

 Information—No action required.
- 4049 Service has paused a delete orphans task to %1 (%2) for Replication Set %3, ID: %4

 Information—No action required.
- 4050 Service has resumed a delete orphans task to %1 (%2) for Replication Set %3, ID: %4
 Information—No action required.
- 4051 Service has stopped a delete orphans task to %1 (%2) for Replication Set %3, ID: %4

 Information—No action required.
- 4052 Service has completed a delete orphans task to %1 (%2) for Replication Set %3, ID: %4
 Information—No action required.
- 4053 Service has started a restore task to %1 (%2) for Replication Set %3, ID: %4

 Information—No action required.

- 4054 Service has paused a restore task to %1 (%2) for Replication Set %3, ID: %4

 Information—No action required.
- 4055 Service has resumed a restore task to %1 (%2) for Replication Set %3, ID: %4
 Information—No action required.
- 4056 Service has stopped a restore task to %1 (%2) for Replication Set %3, ID: %4

 Information—No action required.
- 4057 Service has completed a restore task to %1 (%2) for Replication Set %3, ID: %4
 Information—No action required.
- 4058 Service has started a verification task to %1 (%2) for Replication Set %3, ID: %4

 Information—No action required.
- 4059 Service has paused a verification task to %1 (%2) for Replication Set %3, ID: %4

 Information—No action required.
- 4060 Service has resumed a verification task to %1 (%2) for Replication Set %3, ID: %4

 Information—No action required.
- 4061 Service has stopped a verification task to %1 (%2) for Replication Set %3, ID: %4

 Information—No action required.
- 4062 Service has completed a verification task to %1 (%2) for Replication Set %3, ID: %4

 Information—No action required.
- 4063 Bandwidth limit to %1 (%2) has changed to %3.

Information—No action required.

4100 Product activation code is invalid. Please check that it is typed correctly and is valid for the version of the operating system in use.

Error—If you are in the process of installing Double-Take, verify that you are using a 24 character alpha-numeric code. If Double-Take is already installed, confirm that the code entered is correct. If the code appears to be correct, contact technical support.

4101 This service will not run on this device. Contact your sales representative for upgrade procedures.

Error—The activation code does not match the type of server you are attempting to run on. Contact your vendor for a new activation code or contact technical support.

4110 Target cannot write %1 due to target disk being full. Operation will be retried (%2 times or forever)

Warning—The disk on the target is full. The operation will be retried according to the TGExecutionRetryLimit setting.

4111 Target can not write %1 due to a sharing violation. Operation will be retried (%2 times or forever)

Warning—A sharing violation error is prohibiting Double-Take from writing on the target. The operation will be retried according to the TGExecutionRetryLimit setting.

4112 Target can not write %1 due to access denied. Operation will be retried (%2 times or forever)

Warning—An access denied error is prohibiting Double-Take from writing on the target. The operation will be retried according to the TGExecutionRetryLimit setting.

4113 Target can not write %1 due to an unknown reason. Operation will be retried (%2 times or forever). Please check the log files for further information on the error.

Warning—An unknown error is prohibiting Double-Take from writing on the target. The operation will be retried according to the TGExecutionRetryLimit setting.

4120 Target write to %1 was completed successfully after %2 retries.

Information—No action required.

4150 Target write %1 failed after %2 retries and will be discarded. See the event log or log files for error conditions. After correcting the problem, you should re-mirror or run a verify to resynchronize the changes.

Error—The operation has been retried according to the TGExecutionRetryLimit setting but was not able to be written to the target and the operation was discarded. Correct the problem and remirror the files.

4200 In band task %1 submitted from %2 by %3 at %4

Information—No action required.

4201 In band task %1 discarded (submitted from %2 by %3 at %4)

Warning—A task may be discarded in the following scenarios: all connections to a target are manually disconnected, replication is stopped for all connections to a target, or an auto-disconnect occurs. If one of these scenarios did not cause the task to be discarded, contact technical support.

4202 Running %1 in band script: %2 (task %3 submitted from %4 by %5 at %6)

Information—No action required.

4203 Completed run of in band script: %1 (exit code %2)

Information—No action required.

4204 Error running in band script: %1

Error—Review the task and its associated script(s) for syntax errors.

4205 Timeout (%1 seconds) running in band script: %2

Warning—The timeout specified for the script to complete has expired. Normal processing will continue. You may need to manually terminate the script if it will never

complete.

4206 Run timeout disabled for in band script: %1

Warning—The timeout period was set to zero (0). Double-Take will not wait for the script to complete before continuing. No action is required.

4207 In band scripts disabled by server - no attempt will be made to run %1

Warning—Enable task command processing.

4300 A connection request was received on the target before the persistent target paths could be loaded.

Error—You may need to disconnect and reconnect your replication set.

4301 Unable to block target paths, the driver is unavailable.

Error—If you need to block your target paths, contact technical support.

4302 Target Path %1 has been successfully blocked

Information—No action required.

4303 Blocking of target path: %1 failed. Error Code: %2

Warning—If you need to block your target paths, contact technical support.

4304 Target Path %1 has been successfully unblocked

Information—No action required.

4305 Unblocking of target path: %1 failed. Error Code: %2

Warning—If you need to unblock your target paths, contact technical support.

4306 Target paths for source %1 (%2) Connection id: %3 are already blocked

Warning—No action required.

4307 Target paths for source %1 (%2) Connection id: %3 are already unblocked

Warning—No action required.

4308 Error loading target paths for blocking, registry key %1 has been corrupted.

Error—If you need to block your target paths, contact technical support.

5000 Server Monitor service was successfully started.

Information—No action required.

5001 Server Monitor service was successfully stopped.

Information—No action required.

5002 Placeholders were modified to %1.

Information—No action required.

5100 Failover completed for %1.

Information—No action required.

5101 IP address %1 with subnet mask %2 was added to target machine's %3 adapter.

Information—No action required.

5102 %1 has reached a failover condition. A response from the user is required before failover can take place.

Warning—User intervention has been configured. Open the Failover Control Center and accept or decline the failover prompt.

5103 Started adding drive shares from %1 to %2.

Information—No action required.

5104 %1 drive shares were taken over by %2.

Information—No action required.

5105 Attempting to run the %1 script.

Information—No action required.

5106 The %1 script ran successfully.

Information—No action required.

5107 Error occurred in running %1 script.

Error—Verify that the script identified exists with the proper permissions.

5108 The source machine %1 is not responding to a ping.

Error—This occurs when all monitored IP addresses on the source machine stop responding to pings. Countdown to failover will begin at the first occurrence and will continue until the source machine responds or until failover occurs.

5109 The public NIC on source machine %1 is not responding to a ping.

Error—The failover target did not receive an answer to its ping of the source machine. Eventually, a failover will result. Investigate possible errors (down server, network error, etc.).

5200 Failback completed for %1.

Information—No action required.

5201 IP address %1 was removed from target machine's %2 adapter.

Information—No action required.

5202 Unable to Failback properly because IP address %1 was missing a corresponding SubNet Mask.

Error—Contact technical support.

5300 The following IP address was added to target's monitoring list: %1

Information—No action required.

5301 The following IP address was removed from target's monitoring list: %1

Information—No action required.

5302 Drive share information for %1 has been updated on the target machine.

Information—No action required.

5400 Broadcasted new MAC address %1 for IP address %2.

Information—No action required.

5500 Could not connect to e-mail server. Check to make sure the SMTP server %1 is available (error code: %2).

Warning—Double-Take could not connect to your SMTP server or the username and/or password supplied is incorrect. Verify that SMTP server is available and that you have identified it correctly in your e-mail notification configuration. Also verify that your username and password have been entered correctly.

5501 E-mail notification could not be enabled (error code: %1).

Warning—This alert occurs if there is an unexpected error enabling e-mail notification during service startup. Check to see if any other errors related to e-mail notification have been logged. Also, check to make sure the Windows Management Instrumentation (WMI) service is enabled. If neither of these apply, contact technical support.

5502 E-mail notification could not be initialized. Check to make sure Internet Explorer 5.0 or later is installed.

Warning—E-mail notification no longer requires Internet Explorer 5.0 or later. If you receive this error, contact technical support.

5503 E-mail notification could not be processed. Check to make sure the correct version of SMTPMail.DLL is registered on the system (error code: %1).

Warning—If you are using Double-Take 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 Could not load LocalRS.dll (for e-mail notification).

Warning—This alert occurs if there is an error loading the resource DLL for the service. Typically, this is caused by a missing LocalRS.dll file. Reinstall the software, using the installation Repair option, to install a new copy of the LocalRS.dll. Contact technical support if this error persists.

5505 E-mail could not be sent. Check e-mail settings (error code: %1).

Warning—Verify that the e-mail server that you have identified in your e-mail notification configuration is correct.

5506 One or more required e-mail settings have not been specified (error code: %1).

Warning—At a minimum, you must specify the e-mail server, the From and To addresses, and at least one type of event to include.

5507 E-mail notification could not be initialized. Check to make sure WMI is installed and available (error code: %1).

Warning—If you are using Double-Take 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 An error occurred connecting to the WMI namespace. Check to make sure the Windows Management Instrumentation service is not disabled (error code %1).

Warning—This alert occurs if there is an error with the Windows Management Instrumentation (WMI) service. Verify that you have it installed on the Double-Take server and that it is enabled.

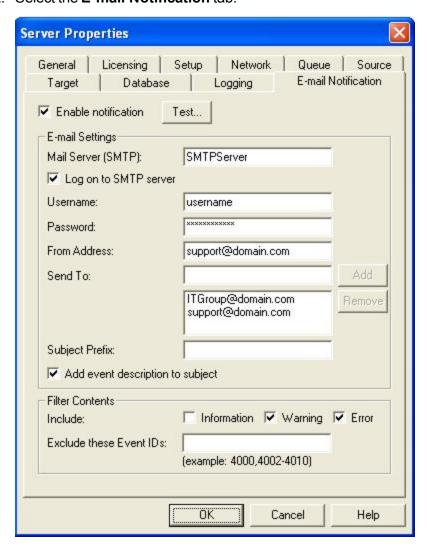
5600 Part or all of the e-mail setting %1 is not in a valid format.

Warning—Verify that the include categories and exclude ID list are identified and formatted correctly.

E-mailing system messages

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

- 1. To enable e-mail notification for a server, right-click the server in the left pane of the Replication Console and select **Properties**.
- 2. Select the E-mail Notification tab.



3. Select Enable notification.



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

4. Specify your e-mail settings.

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



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

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

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

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

If desired, enable **Add event description** to subject to have the description of the message appended to the end of the subject line. This field is optional.

• Filter Contents—Specify which messages that you want to be sent via e-mail. Specify Information, Warning, and/or Error. You can also specify which messages to exclude based on the message ID. Enter the message IDs as a comma or semicolon separated list. You can indicate ranges within the list.



You can test e-mail notification by specifying the options on the E-mail Notification tab and clicking **Test**. If desired, you can send the test message to a different e-mail address by selecting **Send To** and entering a comma or semicolon separated list of addresses. Modify the message text up to 1024 characters, if necessary. Click **Send** to test the e-mail notification. The results will be displayed in a message box.



Click **OK to** close the message and click **Close** to return to the E-mail Notification tab.

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

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

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

Statistics

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

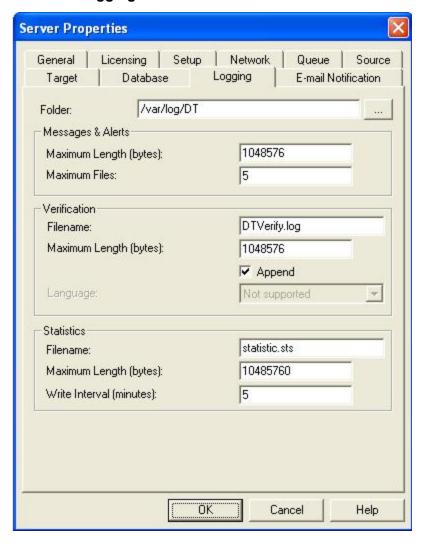
The statistics log file created is a binary file. To view the log file, you must run the DTStat utility from the command prompt.

Sample DTStat output

```
0/11/10 12:48:05:2040
SYSTEMALLOCATOR::Total Bytes: 0
IQALLOCATOR::Total Bytes: 0
SECURITY::Logins : 1 FailedLogins : 0
KERNEL::SourceState: 2 TargetState: 1 Start Time: Tue Sep 11 12:45:26 2007
RepOpsGenerated: 436845 RepBytesGenerated: 0
MirOpsGenerated: 3316423 MirBytesGenerated: 108352749214952
   FailedMirrorCount: 0 FailedRepCount: 0
   ActFailCount: 0 TargetOpenHandles: 0 DriverQueuePercent: 0
TARGET:: PeerAddress: 10.10.1.104 LocalAddress: 10.10.1.104
   Ops Received: 25 Mirror Ops Received: 23
   Retries: 0 OpsDropped: 0 Ops Remaining: 0
Orphan Files Removed: 0 Orphan Directories Removed: 0 Orphan Bytes Removed: 0
   Bytes In Target Queue: 0 Bytes In Target Disk Queue: 0 TasksSucceeded: 0 TasksFailed: 0 TasksIgnored: 0
SOURCE::autoDisConnects: 0 autoReConnects: 1
   lastFileTouched : /log/data file
CONNECTION:: conPeerAddress: 10.10.1.104
   connectTime: Tue Sep 11 12:45:34 2007
   conState: 1 conOpsInCmdQueue: 0 conOpsInAckQueue: 0
   conOpsInRepQueue: 0 conOpsInMirQueue: 0 conBytesInRepQueue: 0
   conOpsTx: 27 conBytesInMirQueue: 0 conBytesTx: 14952687269
   conBytesCompressedTx: 14952
   conOpsRx: 201127 conBytesRx: 647062280 conResentOpCount: 0 conBytesInDiskQueue: 0
   conBandwidthLimit: 429496295 conBytesSkipped: 22867624 conMirrorBytesRemain: 0
   conMirrorPercent: 100.0%
   conTaskCmdsSubmitted: 0 conTaskCmdsQueued: 0
   conTasksSucceeded: 0 conTasksFailed: 0 conTasksIgnored: 0
```

Configuring the properties of the statistics file

- 1. Right-click a machine in the left pane of the Replication Console and select **Properties**.
- 2. Select the **Logging** tab.



- 3. At the top of the tab, specify the **Folder** where the log files for messages, alerts, verification, and statistics will be saved.
- 4. Under **Statistics**, specify the following information.
 - Filename—The name of the statistics log file. The default file name is statistic.sts.
 - Maximum Length—The maximum length of the statistics log file. The default maximum length is 10 MB. Once this maximum has been reached, Double-Take 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. Select the **Setup** tab.
- 6. Verify that Log Statistics Automatically is enabled. If disabled, statistics will not be logged.
- 7. Click **OK** to save the settings.

Viewing the statistics file

The statistics log file created is a binary file. To view the log file, you must run the DTStat utility from a command prompt. From the directory where Double-Take is installed, run the DTStat command.

Command

DTSTAT

Description

Starts the DTStats statistics logging utility from a command prompt

Syntax

DTSTAT [-p][-i < interval>][-t < filename>] [-f < filename>] [-s < filename>] [-st < filename>][-IP < address>] [-START < mm/dd/yyyy hh:mm>][-STOP < mm/dd/yyyy hh:mm>] [-SERVER < ip_address> < port_number>]

Options

- -p—Do not print the output to the screen
- -i *interval*—Refresh from shared memory every interval seconds
- -t filename—Save the data from memory to the specified binary file filename
- -f filename—Reads from a previously saved binary file, filename, that was generated using the -t option instead of reading from memory
- -s *filename*—Saves only the connection data from the data in memory to an ASCII, comma-delimited file, filename
- -st filename—Saves only the target data from the data in memory to an ASCII, comma-delimited file, filename
- -f *filename1* -s *filename2*—Saves only the connection data from a previously saved binary file, filename1, to an ASCII, comma-delimited file, filename2
- -f *filename1* -st *filename2*—Saves only the target data from a previously saved binary file, filename1, to an ASCII, comma-delimited file, filename2
- -IP address—Filters out the specified address in the IP address field and prints only those entries. Specify more than one IP address by separating them by a comma.
- -START mm/dd/yyyy hh:mm—Filters out any data prior to the specified date and time
- -STOP mm/dd/yyyy hh:mm—Filters out any data after the specified date and time
- -SERVER ip_address port_number—Connects DTStat to the specified IP address using the specified port number instead of to the local machine

Examples

- DTStat -i 300
- DTStat -p -i 300 -t AlphaStats.sts
- DTStat -f AlphaStats.sts -s AlphaStats.csv -start 02/02/2007 09:25
- DTStat -server 206.31.4.51 1106

Notes

- This command is not case-sensitive.
- If no options are specified, DTStat will print the output to the screen at an interval of every one second.
- If the statistics are not changing, DTStat will discontinue writing until statistics begin updating again.

Statistics

The following table identifies the Double-Take statistics.



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

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

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

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

Date/Time Stamp

The date and time that the snapshot was taken. This is the date and time that each statistic was logged. By default, these are generated once a second, as long as there are statistics being generated. If mirroring/replication is idle, then DTStat will be idle as well.

System Allocator, Total Bytes

The number of bytes currently allocated to the system pagefile

IQAllocator, Total Bytes

The number of bytes currently allocated to the intermediate queue

Security, Logins

The number of successful login attempts

Security, Failed Logins

The number of failed login attempts

Kernel, SourceState

- 0—Source is not running
- 1—Source is running without the replication driver
- 2—Source is running with the replication driver

Kernel, TargetState

- 0—Target is not running
- 1—Target is running

Kernel, Start Time

Date and time stamp indicating when the Double-Take service was loaded

Kernel, RepOpsGenerated

The number of replication operations generated by the file system driver. An op is a file system operation. Double-Take replicates data by sending the file system operations across the network to the target. RepOpsGenerated indicates the number of file system operations that have been generated by replication.

Kernel, RepBytesGenerated

The number of replication bytes generated by the file system driver. This is the number of bytes generated during replication. In other words, this is roughly the amount of traffic being sent across the network that is generated by replication. It does not take into account TCP/IP overhead (headers and such).

Kernel, MirOpsGenerated

The number of mirror operations transmitted to the target. Mirroring is completed by transmitting the file system operations necessary to generate the files on the target. This statistic indicates the number of file system operations that were transmitted during the initial mirror. It will continue to increase until the mirror is complete. Any subsequent remirrors will reset this field to zero and increment from there.

Kernel, MirBytesGenerated

The number of mirror bytes transmitted to the target. This is the number of bytes generated during mirroring. In other words, this is roughly the amount of traffic being sent across the network that is generated by the mirror. It does not take into account TCP/IP overhead (headers and such). Again, any subsequent remirror will reset this field to zero and increment from there.

Kernel, FailedMirrorCount

The number of mirror operations that failed due to an error reading the file from the disk

Kernel, FailedRepCount

The number of replication operations that failed due to an error reading the file from the disk

Kernel, ActFailCount

The number of activation code failures when loading the source or target. Activation codes can be bad for reasons such as: expiration of evaluation codes, duplicate codes, incorrect codes, etc.

Kernel, TargetOpenHandles

The number of handles currently open on the target

Kernel, DriverQueuePercent

The amount of throttling calculated as a percentage of the stop replicating limit

Target, PeerAddress

The IP address of the source machine

Target, LocalAddress

The IP address of the target machine.

Target, Ops Received

The total number of operations received by this machine as a target since the Double-Take service was loaded

Target, Mirror Ops Received

The total number of mirror operations received by this machine as a target since the Double-Take service was loaded. This number does not reset to zero for remirrors.

Target, Retries

The number of retries performed before all operations were completed

Target, OpsDropped

The number of operations skipped during a difference mirror. During a difference mirror, if Double-Take detects that there have been no changes to a file, then it will indicate the number of operations it did not send for this file in this field.

Target, Ops Remaining

The total number of operations that are left in the target queue

Target, Orphan Files Removed

The number of orphan files removed from the target machine

Target, Orphan Directories Removed

The number of orphan directories removed from the target machine

Target, Orphan Bytes Removed

The number of orphan bytes removed from the target machine

Target, Bytes In Target Queue

The number of bytes currently in the system memory queue on the target

Target. Bytes In Target Disk Queue

The number of bytes currently in the disk queue on the target

Target, TasksSucceeded

The number of task commands that have succeeded on the target

Target, TasksFailed

The number of task commands that have failed on the target

Target, TasksIgnored

The number of task commands that have been ignored on the target

Source, autoDisConnects

The number of automatic disconnects since starting Double-Take. Auto-disconnects occur because the source no longer sees the target This could be because the connection between the two has failed at some point or because the target machine data is changing on the source faster than the source can get the data to the target. This field tracks the number of times an auto-disconnect has occurred since the Double-Take service was started.

Source, autoReConnects

The number of automatic reconnects since starting Double-Take. Auto-reconnect occurs after a target machine is back online. This field tracks the number of times an auto-reconnect has happened since the Double-Take service was started.

Source, lastFileTouched

The last filename that had a replication operation executed

Connection, conPeerAddress

The IP address of the target machine

Connection, connectTime

The time that this connection was established

Connection, conState

The state of the active connection

- 0—None. This indicates a connection has not been established. Statistics are still available for the source and target machines.
- 1—Active. This indicates that the connection is functioning normally and has no scheduling restrictions imposed on it at this time. (There may be restrictions, but it is currently in a state that allows it to transmit.)
- 2—Paused. This indicates a connection that has been paused.
- 4—Scheduled. This indicates a connection that is not currently transmitting due to scheduling restrictions (bandwidth limitations, time frame limitations, and so on).
- 8—Error. This indicates a connection that is not transmitting because something has gone wrong (for example, lost connection).

Only the Scheduled and Error states can coexist. All other states are mutually exclusive. Statistics will display a conState of 12 when the connection is in both a scheduled and an error state because this is the sum of the two values (4 + 8).

Connection, conOpsInCmdQueue

The number of operations waiting to be executed on the target

Connection, conOpsInAckQueue

The number of operations waiting in the acknowledgement queue. Each operation that is generated receives an acknowledgement from the target after that operation has been received by the target. This statistic indicates the number of operations that have yet to receive acknowledgement of receipt.

Connection, conOpsInRepQueue

The number of replication operations currently waiting to be executed on the target

Connection, conOpsInMirQueue

The number of mirror operations currently waiting to be executed on the target

Connection, conBytesInRepQueue

The number of replication bytes remaining to be transmitted to the target

Connection, conOpsTx

The number of operations transmitted to the target. This is the total number of operations that Double-Take has transmitted as a source. In other words, the cumulative number of operations transmitted by this source to all connected targets.

Connection, conBytesInMirQueue

The number of mirror bytes remaining to be transmitted to the target

Connection, conBytesTx

The number of bytes transmitted to the target. This is the total number of bytes that Double-Take has transmitted as a source. In other words, the cumulative number of bytes transmitted by this source to all connected targets.

Connection, conBytesCompressedTx

The number of compressed bytes transmitted to the target.

Connection, conOpsRx

The number of operations received by the target. The number of operations that the target for this connection (as indicated by the IP address field) has received from this source.

Connection, conBytesRx

The number of bytes received by the target. The number of bytes that the target for this connection (as indicated by the IP address field) has received from this source.

Connection, conResentOpCount

The number of operations resent because they were not acknowledged

Connection, conBytesInDiskQueue

The number of bytes in the source disk queue

Connection, conBandwidthLimit

The amount of bandwidth that may be used to transfer data

Connection, conBytesSkipped

The number of bytes skipped during a difference mirror. During a difference mirror, if Double-Take detects that there have been no changes to a file, then it will indicate the number of bytes it did not send for this file in this field.

Connection, conMirrorBytesRemaining

The number of mirror bytes remaining to be transmitted

Connection, conMirrorPercent

The percentage of the mirror that has been completed. This field is determined if the replication set size was calculated.

Connection, conTaskCmdsSubmitted

The number of task commands that have been submitted on the source

Connection, conTaskCmdsQueued

The number of task commands that have been queued on the source

Connection, conTasksSucceeded

The number of task commands that have succeeded on the source

Connection, conTasksFailed

The number of task commands that have failed on the source

Connection, conTasksIgnored

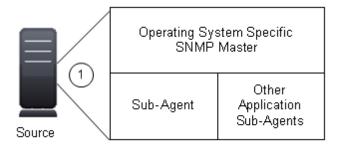
The number of task commands that have been ignored on the source

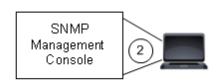
SNMP

SNMP, Simple Network Management Protocol, is the Internet's standard for remote monitoring and management of hosts, routers and other nodes and devices on a network. Double-Take 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.
- 2. 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.







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

Ubuntu SNMP is not supported.

Configuring SNMP on your server

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

- 1. Stop the SNMP daemon (snmpd).
- 2. Make a backup copy of the SNMP configuration file snmpd.conf.
- 3. In order to run the Double-Take 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 &



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 traps

The following table lists the Double-Take SNMP traps.

Kernel, dttrapKernelStarted

Double-Take has started

Kernel, dttrapKernelStopped

Double-Take has stopped

License, dttrapLicenseViolationStartingSource

The source cannot be started due to a license violation

License, dttrapLicenseViolationOnNetwork

A Double-Take serial number conflict was identified on the network

Source, dttrapSourceStarted

Double-Take source component has started

Source, dttrapSourceStopped

Double-Take source component has stopped

Target, dttrapTargetStarted

Double-Take target component has started

Target, dttrapTargetStopped

Double-Take target component has stopped

Connection, dttrapConnectionRequested

The source has requested a connection to the target

Connection, dttrapConnectionRequestReceived

The target has received a connection request from the source

Connection, dttrapConnectionSucceeded

The source to target connection has been established

Connection, dttrapConnectionPause

The source to target connection has paused

Connection, dttrapConnectionResume

The source to target connection has resumed

Connection, dttrapConnectionFailed

The source to target connection was not successful

Connection, dttrapConnectionLost

The source to target connection has been disconnected

Connection, dttrapMemoryLimitReached

The Double-Take memory pool limit has been reached

Connection, dttrapMemoryLimitRemedied

The memory pool usage is below the maximum limit specified

Connection, dttrapAutoReconnect

Auto-reconnect needs to make a new connection

Connection, dttrapScheduledConnectStart

A scheduled connection has been established

Connection, dttrapScheduledConnectEnd

A scheduled end connection has been reached and the connection has been disconnected

Connection, dttrapAutoDisconnectWriteQueue

Auto-disconnect has forced the queue to be written to disk

Connection, dttrapAutoDisconnectPauseTransmission

Auto-disconnect requested that the source pause any operation (create, modify, or delete) sending

Connection, dttrapAutoDisconnectEndConnection

Auto-disconnect has intentionally dropped the connection

Connection, dttrapAutoDisconnectShutdown

Auto-disconnect forced Double-Take to shutdown

Replication, dttrapReplicationStart

Replication has started

Replication, dttrapReplicationStop

Replication has stopped

Mirroring, dttrapMirrorStart

Mirroring has started

Mirroring, dttrapMirrorStop

Mirroring has stopped

Mirroring, dttrapMirrorPause

Mirroring has paused

Mirroring, dttrapMirrorResume

Mirroring has resumed

Mirroring, dttrapMirrorEnd

Mirroring has ended

Verification, dttrapVerificationStart

Verification has started

Verification, dttrapVerificationEnd

Verification has ended

Verification, dttrapVerificationFailure

Verification has failed

Restoration, dttrapRestoreStarted

Restoration has started

Restoration, dttrapRestoreComplete

Restoration is complete

Replication Sets, dttrapRepSetModified

Replication has been modified

Failover, dttrapFailoverConditionMet

Manual intervention is required because failover has detected a failed source machine

Failover, dttrapFailoverInProgress

Failover is occurring

Failover, dttrapTargetFull

The target is full

SNMP statistics

The following table lists the Double-Take SNMP statistics.

General, dtUpTime

Time in seconds since Double-Take was last started

General, dtCurrentMemoryUsage

Amount of memory allocated from the Double-Take memory pool

General, dtMirOpsGenerated

The number of mirror operations (create, modify, or delete) that have been transmitted by the mirroring process

General, dtMirBytesGenerated

The number of bytes that have been transmitted by the mirroring process

General, dtRepOpsGenerated

The number of operations (create, modify, or delete) that have been transmitted by the replication process

General, dtRepBytesGenerated

The number of bytes that have been transmitted by the replication process

General, dtFailedMirrorCount

The number of operations that failed to mirror because they could not be read on the source

General, dtFailedRepCount

The number of operations that failed to be replicated because they could not be read on the source

General, dtActFailCount

The number of activation code errors

General, dtAutoDisCount

The number of auto-disconnects

General, dtAutoReCount

The number of auto-reconnects

General, dtDriverQueuePercent

The amount of throttling calculated as a percentage of the stop replicating limit

Source, dtSourceState

- 0—Source is not running
- 1—Source is running without the replication driver
- 2—Source is running with the replication driver.

Target, dtTargetState

- 0—Target is not running
- 1—Target is running

Target, dtRetryCount

The number of file operations that have been retried

Target, dtOpsDroppedCount

The number of file operations that have failed and will not be retried

Security, dtLoginCount

The number of successful logins

Security, dtFailedLoginCount

The number of unsuccessful logins

Connection, dtConnectionCount

The number of active connections from the source to a target

Connection, dtconlpAddress

The IP address of the connected machine. If at the source, then the IP address of the target. If at the target, then the IP address of the source.

Connection, dtconConnectTime

The duration of time since the connection was first established

Connection, dtconState

The state of the active connection

- 0—None. This indicates a connection has not been established. Statistics are still available for the source and target machines.
- 1—Active. This indicates that the connection is functioning normally and has no scheduling restrictions imposed on it at this time. (There may be restrictions, but it is currently in a state that allows it to transmit.)
- 2—Paused. This indicates a connection that has been paused.
- 4—Scheduled. This indicates a connection that is not currently transmitting due to scheduling restrictions (bandwidth limitations, time frame limitations, and so on).

8—Error. This indicates a connection that is not transmitting because something has gone wrong (for example, lost connection).

Only the Scheduled and Error states can coexist. All other states are mutually exclusive. SNMP will display a dtconState of 12 when the connection is in both a scheduled and an error state because this is the sum of the two values (4 + 8).

Connection, dtconOpsInCmdQueue

The number of operations (create, modify, or delete) in the retransmit queue on the source

Connection, dtconOpsInAckQueue

The number of operations (create, modify, or delete) waiting for verification acknowledgements from the target

Connection, dtconOpsInRepQueue

The number of replication operations (create, modify, or delete) in the queue

Connection, dtconOpsInMirQueue

The number of mirror operations (create, modify, or delete) in the queue

Connection, dtconBytesInRepQueue

The number of bytes in the replication queue

Connection, dtconBytesInMirQueue

The number of bytes in the mirror queue

Connection, dtconOpsTx

The total number of operations (create, modify, or delete) transmitted to the target

Connection, dtconBytesTx

The total number of bytes transmitted to the target

Connection, dtconBytesCompressedTx

The total number of compressed bytes transmitted to the target

Connection, dtconOpsRx

The total number of operations (create, modify, or delete) received from the target

Connection, dtconBytesRx

The total number of bytes received from the target

Connection, dtconResentOpCount

The number of operations that were resent because of acknowledgement errors

Chapter 7 Connections

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.

- 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.

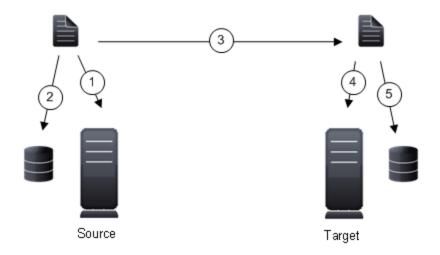
You can perform the following functions to manage your connections.

- Queuing data on page 116
- Reconnecting automatically on page 121
- Pausing and resuming target processing on page 122
- Disconnecting a connection on page 123

Data queues

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.

The following diagram will help you understand how queuing works. Each numbered step is described after the diagram.



- 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.



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

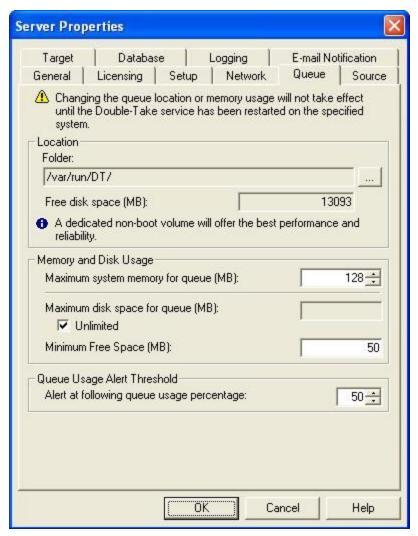
- system memory. As a transaction log is depleted, it is deleted. When all of the transaction log files are deleted, data is again written directly to system memory (step 1).
- 4. To ensure the integrity of the data on the target, the information must be applied in the same order as it was on the source. 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.

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).

Queuing data

You should configure queuing on both the source and target.

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



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

Select a location on a volume that will have minimal impact on the operating system and applications being protected. For best results and reliability, this should be a dedicated, non-boot volume. The disk queue should not be on the same physical or logical volume as the data being replicated.



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

• Maximum system memory for queue—This is the amount of system memory, in MB, that will be used to store data in queues. When exceeded, queuing to disk will be triggered. This value is dependent on the amount of physical memory available but has a minimum of 32 MB. By default, 128 MB of memory is used. If you set it lower, Double-Take Availability will use less system memory, but you will queue to disk sooner which may impact system performance. If you set it higher, Double-Take Availability will maximize system performance by not queuing to disk as soon, but the system may have to swap the memory to disk if the system memory is not available.

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

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

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



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

- 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 Double-Take Availability log. By default, the alert will be generated when the queue reaches 50%.
- 5. Click **OK** to save the settings.

Auto-disconnect and auto-reconnect

While disk queues are user configurable and can be extensive, they are limited. If the amount of disk space specified for disk queuing is met, additional data could not be added to the queue and data would be lost. To avoid any data loss, the auto-disconnect and auto-reconnect processes occur.

- Exhausted queues on the source—If disk queuing is exhausted on the source, Double-Take 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 autoreconnect 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 attempt to persist to disk important target connection information (such as the source and target IP addresses for the connection, various target queue information, the last acknowledged operation, data in memory moved to disk, and so on) before the daemon is stopped. If Double-Take 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.



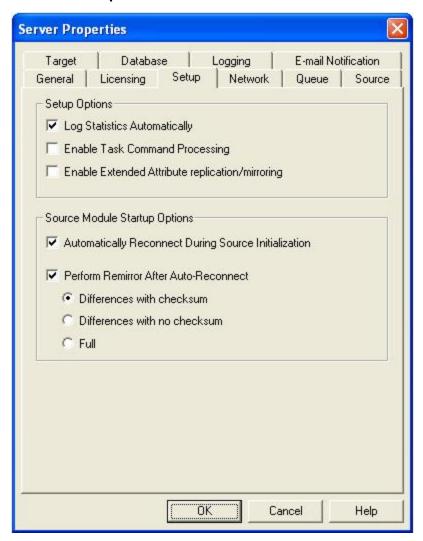
If you are experiencing frequent auto-disconnects, you may want to increase the amount of disk space on the volume where the Double-Take queue is located or move the disk queue to a larger volume. See *Queuing data* on page 116.

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, overwriting the target data changes that you caused, to ensure data integrity between your source and target.

Reconnecting automatically

Use the following steps to configure automatic reconnections.

- 1. Right-click the source server on the left pane of the Replication Console and select **Properties**.
- 2. Select the Setup tab.



- 3. Verify that the check box **Automatically Reconnect During Source Initialization** is marked to enable the auto-reconnect feature.
- 4. Click **OK** to save the settings.

Pausing and resuming target processing

You can break the source/target connection without disconnecting the connection, so that you can control the transmission of data across the network. You can do this by pausing the target. If the target is paused, data is queued on the source until you manually resume the target. For example, you may want to pause the target while you perform a backup of the target data, and then resume the target when the backup is complete.

While the target is paused, the Double-Take source cannot queue data indefinitely. If the source queue is filled, Double-Take will automatically disconnect the connections and attempt to reconnect them. See *Reconnecting automatically* on page 121.

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



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

Disconnecting a connection

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



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

Chapter 8 Mirroring

Mirroring is one of the key components of Double-Take. You can perform the following functions to manage mirroring.

- Stopping, starting, pausing, or resuming mirroring on page 125
- Mirroring automatically on page 127
- Removing orphan files on page 129

Stopping, starting, pausing, or resuming mirroring

After a connection is established, you need to be able to control the mirroring. You can start, stop, pause and resume mirroring. Right-click the connection on the right pane of the Replication Console and select **Mirroring** and the appropriate mirror control.

- Pause or Resume—When pausing a mirror, Double-Take 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.
 - **File differences**—Only those files that are different based size or date and time will be sent from the source to the target. Expand *File difference mirror options compared* below to see how the file difference mirror settings work together, as well as how they work with the global checksum setting on the **Source** tab of the Server Properties.
 - Send data only if Source is newer than Target—Only those files that are newer on the source are sent to the target.



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

- **Use block checksum**—For those files flagged as different, the mirror performs a checksum comparison and only sends those blocks that are different.
- Calculate Replication Set size prior to mirror—Determines the size of the replication set prior to starting the mirror. The mirroring status will update the percentage complete if the replication set size is calculated.

File difference mirror options compared

- **File Differences**—Any file that is different on the source and target based on the date, time, and/or size is transmitted to the target. The mirror sends the entire file.
- File Differences and Only if Source is Newer—Any file that is newer on the source than on the target based on date and/or time is transmitted to the target. The mirror sends the entire file.
- **File Differences and Checksum**—This option is dependent on the global checksum all option on the Server Properties source tab.
 - Checksum All disabled— Any file that is different on the source and target based on date, time, and/or size is flagged as different. The mirror then performs a checksum comparison on the flagged files and only sends those blocks that are different.
 - **Checksum All enabled**—The mirror performs a checksum comparison on all files and only sends those blocks that are different.
- File Differences, Only if Source is Newer, and Checksum—Any file that is newer on the source than on the target based on date and/or time is flagged as different. The mirror then performs a checksum comparison on the flagged files and only sends those blocks that are different.

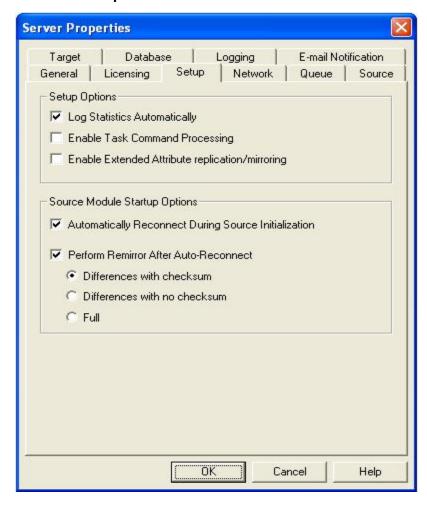
Mirroring automatically

In certain circumstances, for example if the disk-based queues on the source are exhausted, Double-Take 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.



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

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



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



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

- 4. Specify the type of mirror that you wish to perform.
 - **Differences with Checksum**—Any file that is different on the source and target based on date, time, and/or size is flagged as different. The mirror then performs a checksum comparison on the flagged files and only sends those blocks that are different.
 - **Differences with no Checksum**—Any file that is different on the source and target based on date, time, and/or size is sent to the target.
 - Full—All files are sent to the target.



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

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

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

Removing orphan files

An orphan file is a file that exists in the target's copy of the replication set data, but it does not exist in the source replication set data. An orphan file can be created when you delete a file contained in the source replication set while there is no Double-Take 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.

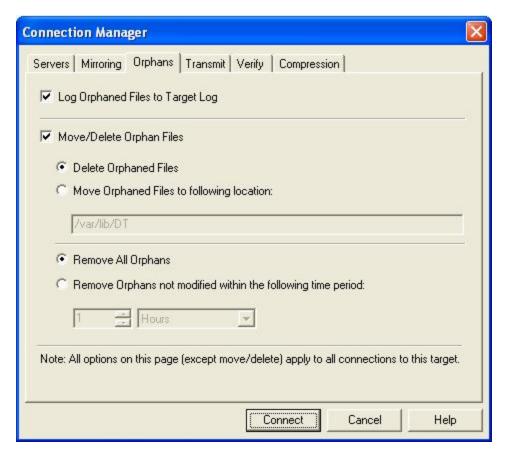


Orphan file configuration is a per target option. All connections to the same target will have the same orphan file configuration.

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

- 1. If you want to preview which files are identified as orphan files, right-click an established connection and select **Remove Orphans**, **Preview**. Check the log file on the target for the list of orphaned files.
- 2. If you want to remove orphan files manually, right-click an established connection and select **Remove Orphans**, **Start**.
- 3. If you want to stop the process after it has been started, right-click the connection and select **Remove Orphans**, **Stop**.
- 4. To configure orphan files for processing during a mirror, verify, or restore, use the following instructions.
 - a. Right-click the connection on the right pane of the Replication Console and select **Connection Manager**.
 - b. Select the **Orphans** tab.



- c. Specify if you want to log the name of the orphan files to the Double-Take log file on the target by marking **Log Orphaned Files to Target Log**.
- d. By default, the orphan files feature is disabled. To enable it, mark **Move/Delete Orphan Files**.
- e. Specify if you want to **Delete Orphaned Files** or **Move Orphaned Files** to a different location. If you select the move option, identify the location where these orphan files will be located.



If you are moving or deleting orphan files, select a move location outside of the replication set. If you select the location where the files are currently located, the files will be deleted. If you select another location inside the replication set, the files will be moved multiple times and then possibly deleted.

- f. Specify if you want to **Remove All Orphans** or **Remove Orphans not modified within the following time period**. If you select the time-based option, only orphans older than the time you specify will be removed.
- g. Click **OK** to save the settings.

Chapter 9 Replication

Replication is one of the key components of Double-Take. This section contains the following replication topics.

- Replication capabilities on page 132—Review this list to learn what Double-Take supports for replication.
- Replication sets on page 134—This section contains instructions for creating and using Double-Take replication sets.
- Starting replication on page 148—Since replication is one of the key components of Double-Take, this topic includes instructions for starting replication.
- Inserting tasks during replication on page 149—You can insert tasks to be processed inline with replication.

Replication capabilities

Double-Take replicates all file and directory data in the supported Linux file systems. Double-Take 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.

- Double-Take is compatible with NFS and Samba services as long as they are mounted on top of Double-Take. Additionally, NFS and Samba should be started after the Double-Take daemon.
- 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. Operations on hard links outside of the replication set are not replicated.
 - If any hard link location existing inside the replication set is deleted on the source, that file or linked location will be deleted from the target.

Replication sets

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 4094 characters. Paths
 longer than 4094 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 Replication Console. You will be notified that the replication set could not be saved. This error checking safeguards the replication set data in the event that more than one client machine is accessing the source's replication sets.
- Double-Take 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.

Creating a replication set

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

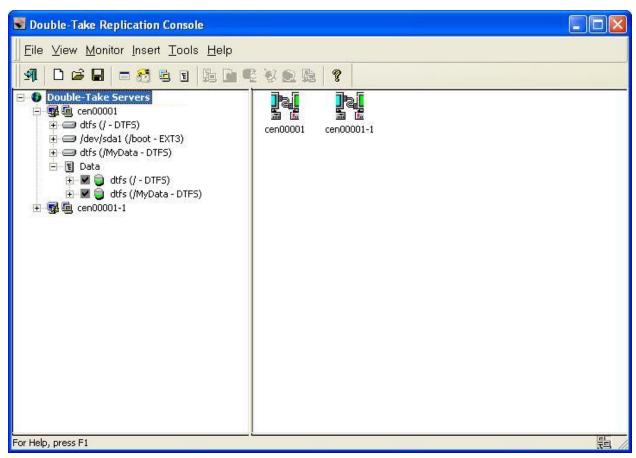
- Highlight a source in the left pane of the Replication Console and select Insert, Replication Set from the menu bar. You can also right-click on the source name and select New, Replication Set.
- 2. A replication set icon appears in the left pane under the source. By default, it is named New Replication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



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

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

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





Be sure and verify what files can be included by reviewing the *Replication capabilities* on page 132.

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

- 5. After selecting the data for this replication set, right-click the new replication set icon and select **Save**. A saved replication set icon will change from red to black.
- 6. If you need to select a block device for replication, right-click the replication set and select **Add Device**.
- 7. The block devices configured for Double-Take Availability replication are shown by default. Highlight the device to include in the replication set and click **OK**.



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



any of these devices, but you cannot use them for Double-Take Availability replication until they are configured for Double-Take Availability replication. The status **no dtloop** indicates the device is not configured for Double-Take Availability replication.

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

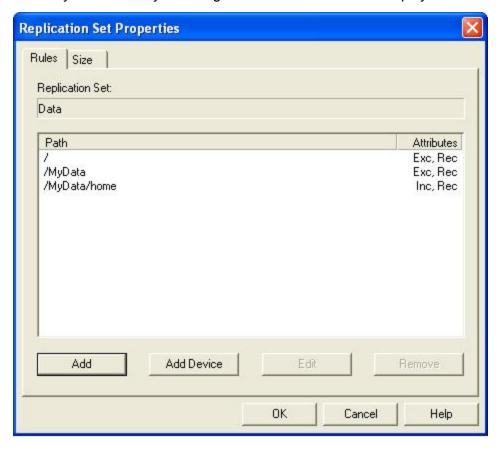
The partition size displayed may not match the output of the Linux df command. This is because df shows the size of the mounted file system not the underlying partition which may be larger. Additionally, Double-Take Availability uses powers of 1024 when computing GB, MB, and so on. The df command typically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

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

Creating or modifying replication rules manually

There may be times when you cannot browse for data when creating a replication set. For example, you can create a replication set rule for a directory or file that does not exist. Since you cannot browse for the location, you have to create replication set rule manually. At other times, the data you want to replicate cannot be easily selected from the Replication Console. For example, you may want to select all .db files from a specific volume or directory. This task may be easier to complete by creating the replication set rule manually. Use the following instructions to create or modify a replication set rule manually.

- 1. If you do not have a replication set created, you need to create one. Highlight a source in the left pane of the Replication Console and select Insert, Replication Set from the menu bar. You can also right-click on the source name and select New, Replication Set. A replication set icon appears in the left pane under the source. By default, it is named New Replication Set. Rename the newly inserted replication set with a unique name by typing over the default name and pressing Enter. This process is similar to naming a new folder in Windows Explorer.
- 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.
- 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**.



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.

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

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

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

Modifying a replication set

Double-Take allows you to make modifications to a replication set when you want to change the data you wish to protect. This allows you to add, remove, or modify any replication set rules without having to create a new replication set.

- 1. In the left pane, highlight the replication set you want to modify and expand the volume and directory levels as needed.
- 2. Modify the items by marking or clearing the volume, drive, directory, or file check boxes. Notice the replication set icon has changed from black to red, indicating changes to the replication set rules.
- 3. After updating the rules for this replication set, right-click the replication set icon and select **Save**. A saved replication set icon will change from red to black.



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

Renaming and copying a replication set

To rename or copy a replication set, click once on a highlighted replication set name to edit the field. Specify a unique name and press **Enter**. The process is similar to renaming a folder in Windows Explorer. If the original replication set has not been saved (red icon), the new name replaces the original name. If the original replication set is saved (black icon), the new name creates a copy of the original replication set.

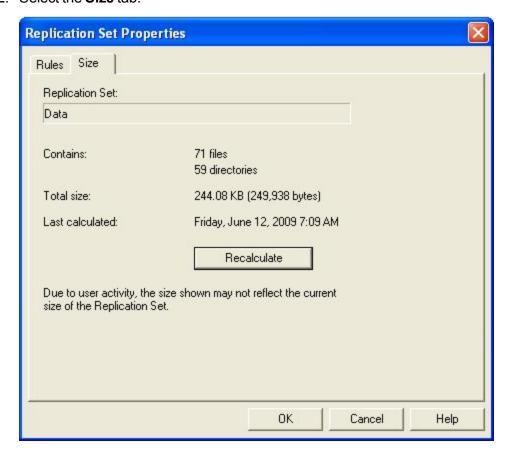


If you save changes to a connected replication set, it is recommended that you perform a mirror to guarantee data integrity between the source and target machines. A dialog box will appear instructing you to disconnect and reconnect the replication set and perform a difference mirror.

Calculating replication set size

While Double-Take is mirroring, the right pane of the Replication Console displays statistics to keep you informed of its progress. If the size of the replication set is determined before the mirror is started, Double-Take 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**.

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



- 3. If the replication set size has never been determined, click **Calculate**. If the replication set has previously been determined, the button will be labeled **Recalculate**. Depending on user activity, the size shown may not accurately reflect the current size of the replication set. If changes are occurring to files in the replication set while the calculation is being made, the actual size may differ slightly. The amount of data is determined at the exact time the calculation is made.
- 4. Click **OK** to return to the Replication Console.



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

If your replication set contains a large number of files, for example, ten thousand or more, you may want to disable the calculation of the replication set size so that data will start

being mirrored sooner. If calculation is enabled, the source calculates the file size before it starts mirroring. This can take a significant amount of time depending on the number of files and system performance. Disabling calculation will result in the mirror status not showing the percentage complete or the number of bytes remaining to be mirrored.

Exporting and importing a replication set

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

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

Deleting a replication set

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

To delete a replication set, right-click the replication set icon and select **Delete**. Additionally, you can highlight the replication set and press the **Delete** key on the keyboard.

Starting replication

Starting replication when establishing a connection is the default and recommended configuration. If replication is not started, data is not added to the queue on the source, and source/target data integrity is not guaranteed.

To start replication, right-click the connection on the right pane of the Replication Console and select **Replication**, **Start**. After starting replication, you should perform a remirror to guarantee the source and target data are identical.

Inserting tasks during replication

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

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

To enable task command processing from the Replication Console, right-click a server in the left pane of the Replication Console, select **Properties**, select the **Setup** tab, and select **Enable Task Command Processing**.



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

Chapter 10 Verification

Verification is the process of confirming that the data on the target is identical to the data on the source. Verification creates a log file detailing what was verified as well as which files are not synchronized. If the data is not the same, Double-Take can automatically initiate a remirror. The remirror ensures data integrity between the source and target.

- Verifying manually on page 151—You can verify your data at any time manually.
- Verifying on a schedule on page 152—You can schedule verification tasks for periodic intervals.
- Configuring the verification log on page 154—You can configure how the verification information is logged.

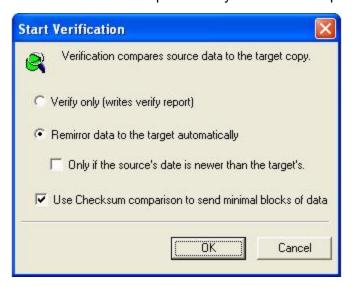


Differences in files on the source and target should be expected for files and applications that are in use during the verification process.

Verifying manually

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

- Right-click the connection on the right pane of the Replication Console and select Verify.
- 2. Select the verification options that you would like to perform.



- **Verify only**—This option verifies the data and generates a verification log, but it does not remirror any files that are different on the source and target.
- Remirror data to the target automatically—This option verifies the data, generates a verification log, and remirrors to the target any files that are different on the source.
- Only if the source's date is newer than the target's—If you are remirroring your files, you can specify that only files that are newer on the source than the target be remirrored.



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

Use Checksum comparison to send minimal blocks of data—Specify if you want the
verification process to use a block checksum comparison to determine which blocks are
different. If this option is enabled, only those blocks (not the entire files) that are different will
be identified in the log and/or remirrored to the target.



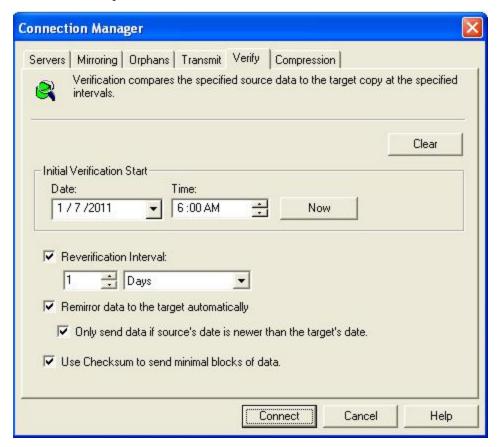
Database applications may update files without changing the date, time, or file size. Therefore, if you are using database applications, you should use the block checksum comparison to ensure proper verification and remirroring.

Click **OK** to start the verification.

Verifying on a schedule

Verification can be scheduled to occur automatically at periodic intervals.

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



- 3. Specify when you want to start the initial verification. Select the immediate date and time by clicking **Now**, or enter a specific **Date** and **Time**. The down arrow next to **Date** displays a calendar allowing easy selection of any date. **Time** is formatted for any AM or PM time.
- 4. Mark the **Reverification Interval** check box to repeat the verification process at the specified interval. Specify an amount of time and choose minutes, hours, or days.
- 5. Select if you want to **Remirror data to the target automatically**. When enabled, Double-Take 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.



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

7. Specify if you want the verification process to **Use Checksum to send minimal blocks of data** to determine which blocks are different. If this option is enabled, only those blocks (not the entire files) that are different will be identified in the log and/or remirrored to the target.



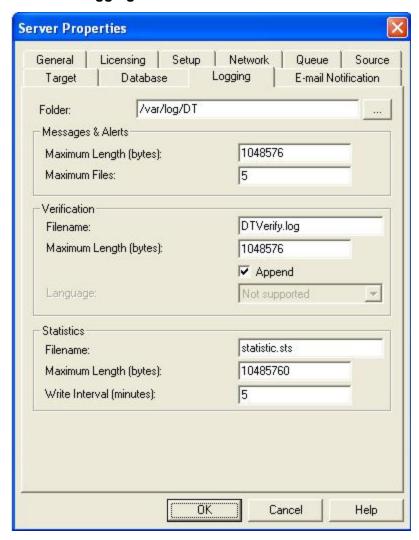
Database applications may update files without changing the date, time, or file size. Therefore, if you are using database applications, you should use the block checksum comparison to ensure proper verification and remirroring.

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

Configuring the verification log

A verification log is created on the source during the verification process. The log identifies what is verified as well as which files are not synchronized.

- 1. Right-click the source server on the left pane of the Replication Console and select **Properties**.
- 2. Select the Logging tab.



- 3. At the top of the window, **Folder** identifies the location where the log files identified on this tab are stored. By default, the log files are stored in the same directory as the Double-Take program files.
- 4. Under the Verification section, **Filename** contains the base log file name for the verification process. The replication set name will be prepended to the base log file name. For example, since the default is DTVerify.log, the verification log for the replication set called UserData would be UserData DTVerify.log.
- 5. Specify the **Maximum Length** of the log file. The default is 1048576 bytes (1 MB). When the log file reaches this limit, no additional data will be logged.
- 6. By default, the log is appended to itself each time a verification process is completed. Clear the

Append check box if you do not want to append to the previous log file.



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

- 7. Specify the **Language** of the log file. Currently, English is the only available language.
- 8. Click **OK** to save the settings.

In the log file, each verification process is delineated by beginning and end markers. A list of files that are different on the source and target is provided as well cumulative totals for the verification process. The information provided for each file is the state of its synchronization between the source and the target at the time the file is verified. If the remirror option is selected so that files that are different are remirrored, the data in the verify log reflects the state of the file before it is remirrored, and does not report the state of the file after it is remirrored. If a file is reported as different, review the output for the file to determine what is different.

Chapter 11 Data transmission

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 transmission to suit your environment.

- Stopping, starting, pausing, and resuming transmission on page 157—You can maintain the
 source/target connection, but still control the transmission of data across the network by using the
 manual transmission controls. If transmission is paused, the data is queued on the source until
 you manually restart the transmission.
- Scheduling data transmission on page 157—You can set event driven or scheduling criteria to
 determine when data is transmitted. Data is queued on the source until the event or schedule is
 met. Also, transmission can be stopped by using these criteria. Scheduled transmission options
 can be toggled on and off, allowing you to enable them only when you need to use them.
- Limiting transmission bandwidth on page 162—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.
- Compressing data for transmission on page 164—You can compress data to reduce the amount
 of bandwidth needed to transmit Double-Take data.

Stopping, starting, pausing, and resuming transmission

To start, pause, or resume the transmission of data from the source to the target, right-click an established connection and select **Transmit** and the appropriate transmission control.

Scheduling data transmission

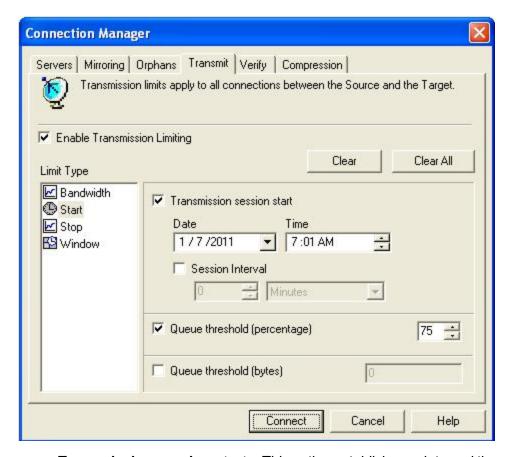
Using the Connection Manager **Transmit** tab, you can set start and stop criteria along with a schedule window.



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

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

- 1. Right-click the connection on the right pane of the Replication Console and select **Connection Manager**.
- Select the Transmit tab. The Transmit tab contains four limit types: Bandwidth, Start, Stop, and Window. The transmission options for each limit type are displayed by highlighting a selection in the Limit Type box.
 - At the top of the **Transmit** tab dialog box, the **Enable Transmission Limiting** check box allows you to turn the transmission options on or off. You can enable the transmission options by marking the **Enable Transmission Limiting** check box when you want the options to be applied, but you can disable the transmission options, without losing the settings, by clearing that check box.
 - Also at the top of the **Transmit** tab dialog box, the **Clear All** button, when selected, will remove all transmission limitations that have been set under any of the limit types. The **Clear** button will clear the settings only for the **Limit Type** selected.
- 3. When you schedule transmission start criteria, transmission will start when the criteria is met and will continue until the queue is empty or a transmission stop criteria is met. Select the **Start option** in the Limit Type box.
 - Define the start options for Double-Take 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

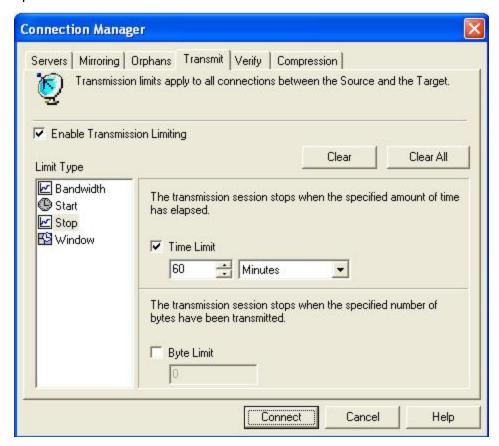
MB. Double-Take will wait until there is 500 MB of data in the queue and then begin transmitting the data. Like other start criteria, Double-Take continues transmitting until the queue is empty or a Double-Take stop criteria is met. Specify a percentage of the disk queue and system memory that must be 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.



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

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

Define the stop options to stop Double-Take 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.



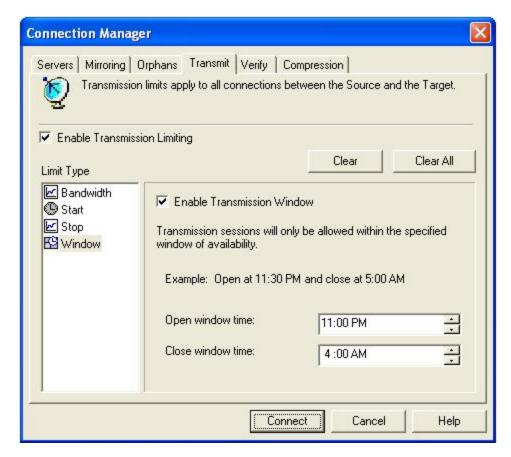
The transmission start and stop criteria should be used in conjunction with each other. For example, if you set the **Queue Threshold** equal to 10 MB and the **Byte Limit** equal to 10 MB, a network connection will be established when there is 10 MB of data in the queue. The data will be transmitted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

5. Schedule a transmission window to establish a period of availability for all Double-Take 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.



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

Limiting transmission bandwidth

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



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

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

- 1. Right-click the connection on the right pane of the Replication Console and select **Connection Manager**.
- 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—Specify the number of kilobits to send every second.



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

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

Compressing data for transmission

To help reduce the amount of bandwidth needed to transmit Double-Take data, compression allows you to compress data prior to transmitting it across the network. In a WAN environment this provides optimal use of your network resources. If compression is enabled, the data is compressed before it is transmitted from the source. When the target receives the compressed data, it decompresses it and then writes it to disk. On a default Double-Take installation, compression is disabled.



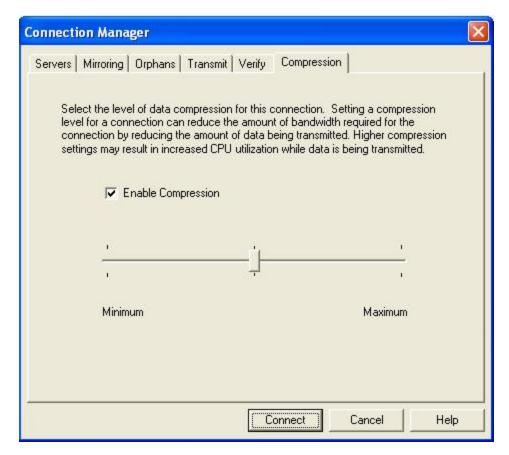
Any replication sets from a source connected to the same IP address on a target will share the same compression configuration.

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

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

Use the following instructions for setting compression.

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



- 3. By default, compression is disabled. To enable it, select **Enable Compression**.
- 4. Depending on the compression algorithms available for your operating system, you may see a slider bar indicating different compression levels. Set the level from minimum to maximum compression to suit your needs.
- 5. Click **OK** to save the settings.

Chapter 12 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 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.

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

- Configuring failover monitoring on page 167
- Editing failover monitoring configuration on page 174
- Monitoring failover monitoring on page 175
- Failing over on page 178
- Removing failover monitoring configuration on page 178

Configuring failover monitoring

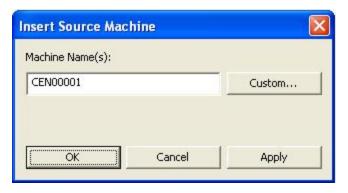
Before beginning your failover configuration, review your IP address and subnet configuration on the source. Because of limitations in the way the Linux kernel handles IP address aliases, you will not be able to mix subnets on the eth0 network interface. Failover should not cause problems in this configuration,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

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



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

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



- Type the name of the machine that you want to monitor in Machine Name(s) and click OK.
- Click **Custom**. Enter the name of the server and click Add. Specify the IP address and subnet mask of the specified server and click **OK**. Click **OK** again.

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

- 6. In the **Names to Monitor** tree, locate and select the IP addresses on the source that you want to monitor.
- 7. Highlight an IP address that you have selected for monitoring and select a **Target Adapter** that

will assume that IP address during failover. Repeat this process for each IP address that is being monitored.

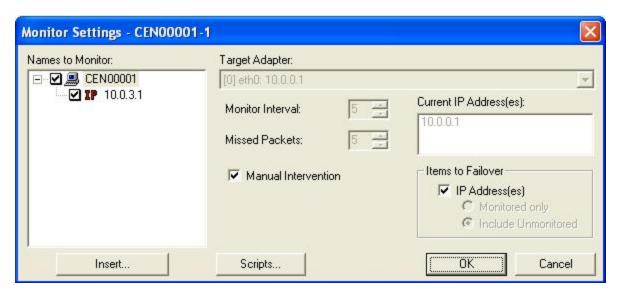


Current IP Addresses displays the IP address(es) currently assigned to the selected target adapter.

- 8. Highlight an IP address that you have selected for monitoring and select a **Monitor Interval**. This setting identifies the number of seconds between the monitor requests sent from the target to the source to determine if the source is online. Repeat this step for each IP address that is being monitored.
- 9. Highlight an IP address that you have selected for monitoring and select the **Missed Packets**. This setting is the number of monitor replies sent from the source to the target that can be missed before assuming the source machine has failed. Repeat this step for each IP address that is being monitored.



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



- 10. Highlight the source name and specify the **Items to Failover**, which identifies which source components you want to failover to the target.
 - **IP Addresses**—If you want to failover the IP addresses on the source, enable this option and then specify the addresses that you want to failover. When the source and target are

on the same subnet, generally a LAN environment, you should failover the IP address. If the source and target are on different subnets, generally a WAN environment, you should not failover the IP address. See *WAN considerations* on page 170 for options on handling WAN failover.

- Monitored only—Only the IP address(es) that are selected for monitoring will be failed over.
- Include Unmonitored—All of the IP address(es) will be failed over.



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

- 11. By default, **Manual Intervention** is enabled, allowing you to control when failover occurs. When a failure occurs, a prompt appears in the Failover Control Center and waits for you to manually initiate the failover process. Disable this option only if you want failover to occur immediately when a failure occurs.
- 12. If you are using any failover or failback scripts, click **Scripts** and enter the path and filename for each script type. Scripts may contain any valid Linux command, executable, or script file. Examples of functions specified in scripts include stopping daemons on the target before failover because they may not be necessary while the target is standing in for the source, stopping daemons on the target that need to be restarted with the source's machine name and IP address, starting daemons or loading applications that are in an idle, standby mode waiting for failover to occur, notifying the administrator before and after failover or failback occurs, stopping daemons on the target after failback because they are no longer needed, stopping daemons on the target that need to be restarted with the target machine's original name and IP address, and so on. Specify each script that you want to run and the following options, if necessary.
- 13. If you want to delay the failover or failback processes until the associated script has completed, mark the appropriate check box.
- 14. If you want the same scripts to be used as the default for future monitor sessions, mark the appropriate check box.
- 15. Click **OK** to return to the Monitor Settings dialog box.
- 16. Click **OK** on the Monitor Settings dialog box to save your monitor settings and begin monitoring for a failure.

WAN considerations

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

- DNS updates—You can script DNS updates to modify, at failover time, the source server's DNS A records to have the IP address of the target. When clients resolve a name to an IP address, they will resolve to the target IP address. Depending on the domain size and how DNS updates are propagated, it may take several minutes or even hours for the updates to complete.
- Reconfigure routers using a failover script—You can automatically reconfigure routers
 using a failover script to move the source's subnet from the source's physical network to the
 target's physical network. Since the route to the source's subnet will be changed at failover, the
 source server must be the only system on that subnet, which in turn requires all server
 communications to pass through a router. Additionally, it may take several minutes or even hours
 for routing tables on other routers throughout the network to converge.
- VPN infrastructure—A VPN infrastructure allows your source and target to be on the same subnet, in which case IP address failover will work the same as a LAN configuration.

Scripting example using the BIND DNS client

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

```
nsupdate "dnsover"
```

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

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

update delete source_server_name.fully_qualified_domain.com

update add target_server_name.fully_qualified_domain.com 86400 A target_server_IP

send
```

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

```
nsupdate "dnsback"
```

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

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

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

Protecting NFS exports

NFS exports must be configured for failover through the failover scripts or created manually on the target after failover.

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

service nfs start

7. If necessary, update DNS.

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

Protecting Samba shares

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

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

service smb start

7. If necessary, update DNS.

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

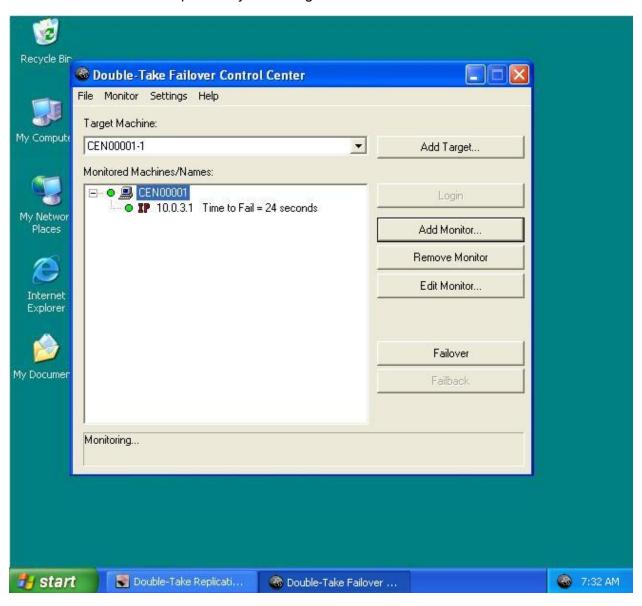
Editing failover monitoring configuration

If you want to edit the monitor settings for a source that is currently being monitored, highlight that source on the Monitored Machines tree on the main Failover Control Center screen and click **Edit**. The Monitor Settings dialog box will open. See *Configuring failover monitoring* on page 167.

Monitoring failover monitoring

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

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





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

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

The following table identifies how the visual indicators change when the source is online.

Time to Fail Countdown

The Time to Fail counter is counting down and resetting each time a response is received from the source machine.

Status Bar

The status bar indicates that the target machine is monitoring the source machine.

Colored Bullets

The bullets are green.

When the Time to Fail value has decreased by 25% of the entire timeout period, the bullet changes from green to yellow, indicating that the target has not received a response from the source. The yellow bullet is a caution signal. If a response from the source is received, the countdown resets and the bullets change back to green. If the countdown reaches zero without the target receiving a response from the source, failover begins.

Desktop Icon Tray

The Windows desktop icon tray contains a failover icon with red and green computers.

The following table identifies how the visual indicators change when the source fails and failover is initiated.

Time to Fail Countdown

The Time to Fail countdown value is 0.

Status Bar

The status bar displays the source machine and IP address currently being assumed by the target.

Colored Bullets

The bullets are red.

Desktop Icon Tray

The Windows desktop icon tray contains a failover icon with red and green computers.

The following table identifies how the visual indicators change when failover is complete.

Time to Fail Countdown

The Time to Fail counter is replaced with a failed message.

Status Bar

The status bar indicates that monitoring has continued.

Colored Bullets

The bullets are red.

Desktop Icon Tray

The Windows desktop icon tray contains a failover icon with a red computer.

Failing over

The failover process, including script processing, can be tested at any time. To force unavailability, disconnect the network cable from a monitored machine, wait for the **Time to Fail** counter to decrease to zero and failover begins. To avoid the countdown delay, highlight the monitored machine name in the Failover Control Center window and select **Failover**.

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



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

When a failure occurs, an alert is forwarded to the Linux system log. You can then start the Failover Control Center and respond to the manual intervention prompt.

If SNMP is installed and configured, an SNMP trap is also generated. When using a third-party SNMP manager, an e-mail or page can be generated to notify you of the failure.

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

Click **Cancel** to abort the failover process. If necessary, you can initiate failover later from the Failover Control Center. Click **OK** to proceed with failover.

Removing failover monitoring configuration

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

Chapter 13 Failback and restoration

Failover occurred because the target was monitoring the source for a failure, and when a failure occurred, the target stood in for the source. User and application requests that were directed to the failed source are routed to the target.

While the users are accessing their data on the target, you can repair the issue(s) on the source. Before users can access the source again, you will need to restore the data from the target back to the source and perform failback. Failback is the process where the target releases the source identity it assumed during failover. Once failback is complete, user and application requests are no longer routed to the target, but back to the source.

Ideally, you want to restore your data from the target back to the source before you failback. This allows users who are currently accessing their data on the target because of failover to continue accessing their data. Restoration before failback reduces user downtime. The other method allows you to failback first and then restore the data from the target to the source. This method may be easier in some situations, but users may experience longer downtime, depending on the amount of data to be restored, because they will be unable to access their data during both the restoration and the failback processes.

- Restoring then failing back on page 180
- Failing back then restoring on page 184

Restoring then failing back

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

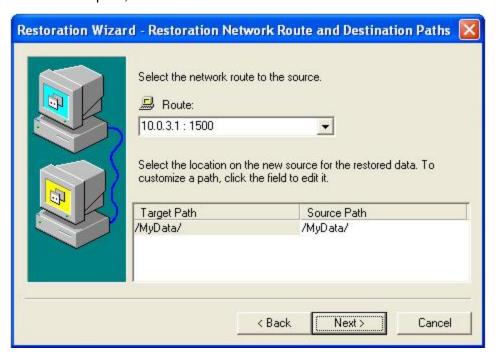
- 1. Resolve the problem(s) on the source that caused it to fail. If you have to rebuild your source, use a unique identity.
- 2. Stop any applications that were failed over that may be running on your source. The files must be closed on the source so that updated files from the target will successfully overwrite the files on the source during the restoration.
- Modify the source so that it can be brought onto the network with a new, unique IP address or one that was not failed over. It needs to be able to exist on the network without an IP address conflict and communicate with the target.
- 4. At this point, confirm you have the following configuration.
 - Your target is standing in for your source because of failover, and users are accessing their data from the target.
 - Your source is back online with a unique IP address.
 - The source and target can communicate with each other.
 - All applications on the source are stopped.
- 5. From your target, confirm the Replication Console is communicating with the source using the new, unique IP address.
 - a. From the Replication Console on the target, right-click the source and select **Remove**.
 - b. Depending on your configuration, the source may be automatically inserted back into the Replication Console. If it is not, select **Insert**, **Server**. Specify the source server by the new IP address and click **OK**.
- 6. Disconnect the connection from the original source to the target, if it still exists.
- 7. Begin your restoration process.
 - a. From the Replication Console, select **Tools**, **Restoration Wizard**.
 - b. Review the Welcome screen and click **Next** to continue.



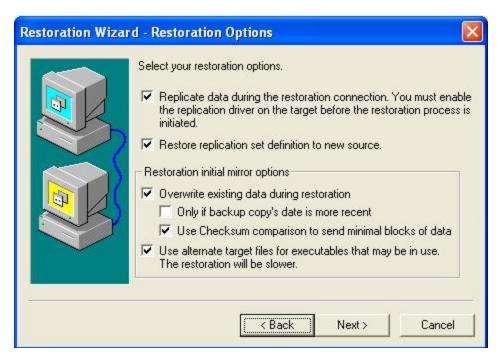
At any time while using the Restoration Wizard, click **Back** to return to previous screens and review your selections.

- c. Select the target that contains the current copy of the data that you want to restore and click **Next**.
- d. Select the original source or **Alternate**, if your original source is not listed. This option identifies to the target which data set you are trying to restore so that the appropriate replication sets can be presented to you.
- e. Click Next to continue.
- f. Specify if you want to use an existing replication set or create a new one. This replication set will be used to connect from the target to the source.
 - **Use this replication set**—If you choose to use an existing replication set, specify the name of that replication set by selecting it from the pull-down menu. You will have an opportunity to modify the replication set definition.

- Create a new replication set with this name—If you choose to create a new replication set, specify a replication set name. With this option, you will need to define the data to be restored.
- g. Click **Next** to continue.
- h. A tree display appears identifying the data available for restoration. Mark the check box of the volumes, directories, and/or files you want to restore. Keep in mind that if you exclude volumes, folders, and/or files that were originally replicated, it may compromise the integrity of your applications or data.
- i. Click Next to continue
- j. Select the new source server. This is the server where the data from the target will be restored. This may be the original source server or a new server. Click **Next** to continue.
- k. Select your network route to the new source, which includes the IP address and port number. Also select the location on the new source for the restored data. If you want to set a customized path, click in the field under **Source Path** to edit the location.



- I. Click **Next** to continue.
- m. Specify the restoration options that you want to use.



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



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

 Use Checksum comparison to send minimal blocks of data—Specify if you want the restoration process to use a block checksum comparison to determine which blocks are different. If this option is enabled, only those blocks (not the entire files) that are different will be restored to the new source.



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

- Use alternate target files for executables that may be in use—If you have executables that may be in use during the restoration, you can have Double-Take create and update an alternate file during the restoration. Once the mirroring and replication operations have been completed, the alternate file will be renamed to the original file. This process will reduce the speed of your restoration, so it should only be used if executables may be in use.
- n. Review your selections and click **Finish** to begin the restoration.
- 8. Monitoring the restoration connection and after the **Mirror Status** is **Idle**, schedule a time for failback. User downtime will begin once failback is started, so select a time that will have minimal disruption on your users.
- 9. When you are ready, begin the failback process.
 - a. Stop user access to the target.
 - b. In the Replication Console, watch the restoration connection until activity has ended and replication is in a **Ready** state. This will happen as the final data in queue, if any, is applied on the source. The replication **Ready** state indicates replication is waiting for new incoming data changes.
 - c. Disconnect the restoration connection.
 - d. Open the Failover Control Center.
 - e. Select the original target that is currently standing in for the original failed source.
 - f. Highlight the failed source and click **Failback**. The user downtime starts now. If you have a pre-failback script configured, it will be started.
 - g. When failback is complete, the post-failback script, if configured, will be started. When the script is complete, you will be prompted to determine if you want to continue failover monitoring, do not select either option. Leave the prompt dialog box open as is.
- 10. On the source, modify the identity back to the original source IP address.
- 11. Confirm the Replication Console is communicating with the source using the original IP address.
 - a. Right-click the source and select **Remove**.
 - b. Depending on your configuration, the source may be automatically inserted back into the Replication Console. If it is not, select **Insert**, **Server**. Specify the source server by the original IP address and click **OK**.
- 12. At this time, you can go back to the dialog box in the Failover Control Center. Select **Continue** or **Stop** to indicate if you want to continue monitoring the source. After you have selected whether or not to continue monitoring the source, the source post-failback script, if configured, will be started.



The source must be online and Double-Take must be running to ensure that the source post-failback script can be started. If the source has not completed its boot process, the command to start the script may be lost and the script will not be initiated.

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

Failing back then restoring

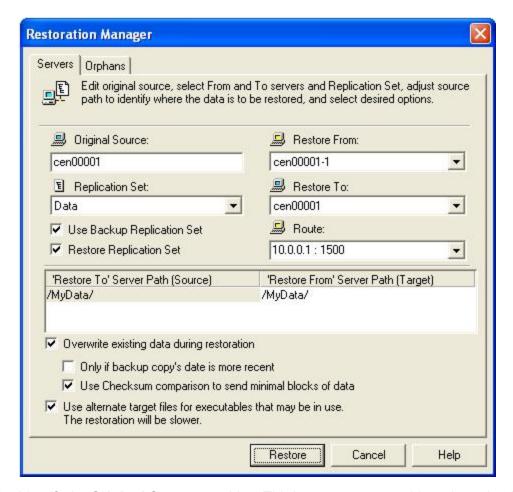
Use these instructions to failback first and then restore your data.

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



The source must be online and Double-Take must be running to ensure that the source post-failback script can be started. If the source has not completed its boot process, the command to start the script may be lost and the script will not be initiated.

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



- 8. Identify the **Original Source** machine. This is your source machine where the data originally resided.
- 9. Select the **Restore From** machine. This is the target machine where the copy of the data is stored.
- 10. Replication Set contains the replication set information stored on the target machine (the machine in Restore From). If no replication sets are available, the list will be blank. Select the replication set that corresponds to the data that you need to restore.
- 11. Select the **Restore To** machine. This is your source where the updated data from the target will be sent.
- 12. Select the Use Backup Replication Set check box to use the target's copy of the replication set database for the restoration. If this check box is not marked, you will be accessing the replication set information from the source.
- 13. Select the **Restore Replication Set** check box to restore the target's copy of the replication set database to the source during the restoration process.
- 14. Select the **Route** on the target. This is the IP address and port on the target that the data will be transmitted through. This allows you to select a different route for Double-Take traffic. For example, you can separate regular network traffic and Double-Take traffic on a machine with multiple IP addresses.
- 15. The **Restore To Server Path** and **Restore From Server Path** paths will automatically be populated when the replication set is selected. The restore to path is the directory that is the

common parent directory for all of the directories in the replication set. If the replication set crosses volumes, then there will be a separate path for each volume. The restore from path is the path on the target server where the replicated files are located.



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

- 16. Select the restoration conditionals that you want to use.
 - Overwrite existing data during restore—This option restores all existing files by overwriting them. Any files that do not exist on the source are written also. If this option is disabled, only files that do not exist on the source will be restored.
 - Only if backup copy's date is more recent—This option restores only those files that are newer on the target than on the source. The entire file is overwritten with this option.



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

- Use Checksum comparison to send minimal blocks of data—Specify if you want the
 restoration process to use a block checksum comparison to determine which blocks are
 different. If this option is enabled, only those blocks (not the entire files) that are different will
 be restored to the source.
- 17. If you want to configure orphan files, click the **Orphans** tab. The same orphan options are available for a restoration connection as a standard connection.
- 18. Click **Restore** to begin the restoration.

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

Chapter 14 Server settings

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

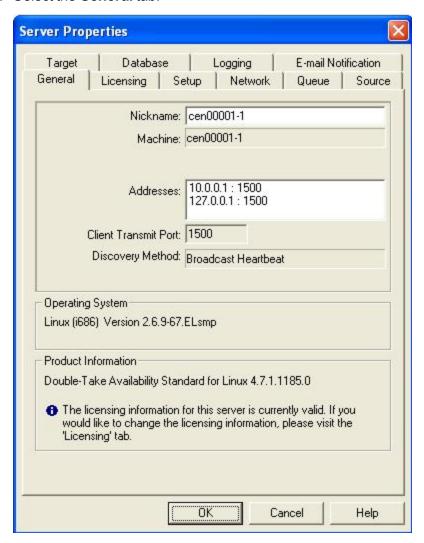
This section contains the following topics, each corresponding to a tab in the Server Properties dialog box.

- Identifying a server on page 188
- Licensing a server on page 190
- Configuring server startup options on page 193
- Configuring network communication properties for a server on page 195
- Queuing data on page 116
- Configuring source data processing options on page 200
- Configuring target data processing options on page 202
- Specifying the Double-Take database storage files on page 203
- Specifying file names for logging and statistics on page 204
- E-mailing system messages on page 92

Identifying a server

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

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



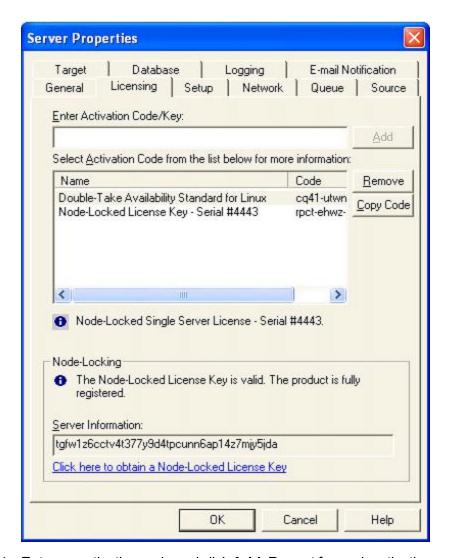
- 4. Specify the server identity information. Some of the fields are informational only.
 - Nickname—A nickname is saved in the Replication Console workspace, therefore, it only
 appears in the Replication Console on this server. It is not communicated across the
 network. If you export a workspace and use it on another Double-Take server, the server
 nickname will appear there also.
 - Machine—This is the actual server name. This field is not modifiable.
 - Addresses—The IP address(es) for this server are listed in this field. This information is not modifiable and is displayed for your information. The machine's primary address is listed first.

- Client Transmit Port—This field displays the port that the Replication Console uses to send commands to a server. This port cannot be modified.
- **Discovery Method**—This field indicates the method in which the Replication Console identifies the Double-Take server.
 - **Manual**—A Double-Take server was manually inserted into the Replication 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.
- 5. Click **OK** to save the settings.

Licensing a server

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

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



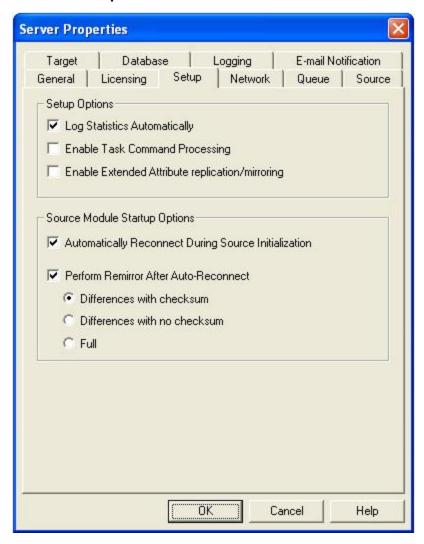
- 4. Enter an activation code and click **Add**. Repeat for each activation code.
- 5. Highlight an activation code in the list to display any status messages for that code below the list display.
- 6. If you need to remove a code from the server, highlight it in the list and click **Remove**.
- 7. To update a temporary node-locked license to a permanent license, you need to provide server information which will be used to generate a permanent node-locked license.
 - a. After entering your temporary node-locked license, click **OK** to activate it. At this point, the temporary license is activated, and you have 14 days to update it to a permanent, nodelocked license.
 - b. Reopen the Server Properties Licensing tab.
 - c. Highlight your temporary node-locked license in the list to display the Node-Locking section at the bottom of the **Licensing** tab.
 - d. Click the hyperlink in the Node-Locking section. If you do not have an Internet connection, copy the **Server Information** text from the Node-Locking section into the form at https://activate.doubletake.com from another machine.
 - e. After you submit the form, you will receive an email with a node-locked license key. Enter that key on the **Licensing** tab and click **Add**. The permanent activation code is specific to

this server. It cannot be used on any other server. If the activation code and server do not match, Double-Take will not run. 8. Click \mathbf{OK} to apply the activation code(s) you entered and/or the node-locked license key.

Configuring server startup options

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

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



- 4. Specify the server setup and source startup options.
 - Log Statistics Automatically—If enabled, Double-Take 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 Replication Console, but it can only be initiated through the scripting language. See the *Scripting Guide* for more information.

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

- Enable Extended Attribute replication/mirroring—This option enables extended
 attribute replication and mirroring. You must enable this option on all of your source and
 target servers.
- Automatically Reconnect During Source Initialization—If enabled, Double-Take 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.

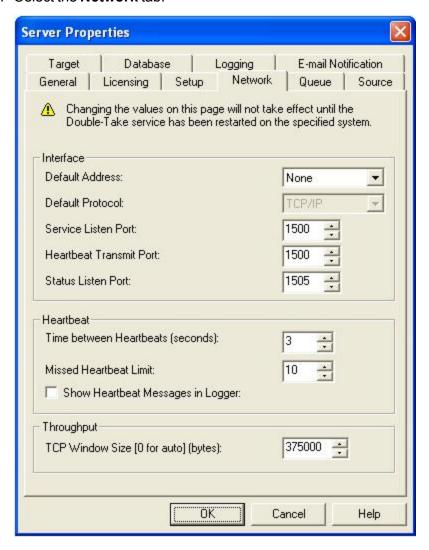


Database applications may update files without changing the date, time, or file size. Therefore, if you are using database applications, you should use the **Differences with checksum** or **Full** option.

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

Configuring network communication properties for a server

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



- 4. Specify the network communication properties.
 - Default Address—On a machine with multiple NICs, you can specify which address
 Double-Take 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 Replication 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 Replication Console. By default, there are 3 seconds between heartbeats. If
 you set this option to 0, the heartbeats are disabled.
- **Missed Heartbeat Limit**—This is the number of heartbeats which can be missed before transmission is stopped and data is queued on the source.
- Show Heartbeat Messages in Logger—This checkbox enables the heartbeat messages in the Double-Take log. Enabling this option will cause your logs to fill up faster.
- TCP Window Size—This option is the size, in bytes, of the buffer used for TCP transfers. This is an operating system buffer, not a Double-Take buffer. If this option is set to zero (0), Linux kernel versions 2.6.7 or later can automatically tune this buffer setting for best server performance. Therefore, the recommended setting is 0 for automatic tuning, if you are using a version 2.6.7 or later Linux kernel. If you want to reduce or control network traffic, you can configure this option to a static size. The default is 375000 for a 1 GB network. Modifications should be relative to that speed using the calculation 37500 * network_ speed_in_bits_per_second / 100 Mbit.



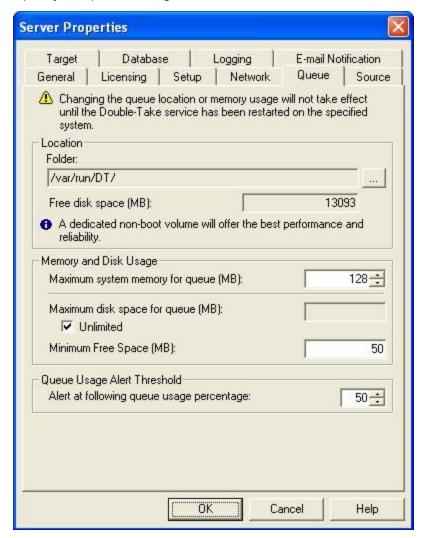
If you want to control network traffic, you may find the Double-Take bandwidth limiting features to be a better method.

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

Queuing data

You should configure queuing on both the source and target.

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



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

Select a location on a volume that will have minimal impact on the operating system and applications being protected. For best results and reliability, this should be a dedicated, non-boot volume. The disk queue should not be on the same physical or logical volume as the data being replicated.



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

• Maximum system memory for queue—This is the amount of system memory, in MB, that will be used to store data in queues. When exceeded, queuing to disk will be triggered. This value is dependent on the amount of physical memory available but has a minimum of 32 MB. By default, 128 MB of memory is used. If you set it lower, Double-Take Availability will use less system memory, but you will queue to disk sooner which may impact system performance. If you set it higher, Double-Take Availability will maximize system performance by not queuing to disk as soon, but the system may have to swap the memory to disk if the system memory is not available.

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

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

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

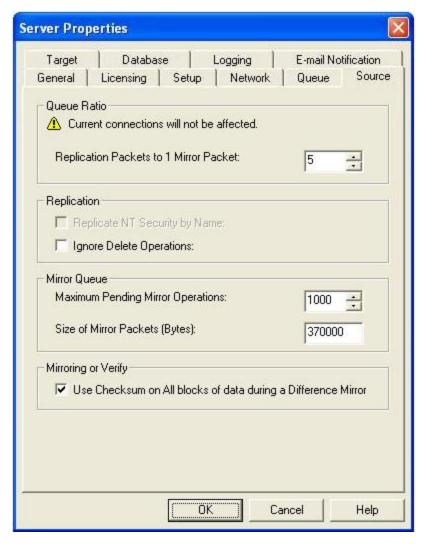


The **Maximum disk space for queue** and **Minimum Free Space** settings work in conjunction with each other. For example, assume your queues are stored on a 10 GB disk with the **Maximum disk space** for queue set to 10 GB and the **Minimum Free Space** set to 500 MB. If another program uses 5 GB, Double-Take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

- 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 Double-Take Availability log. By default, the alert will be generated when the queue reaches 50%.
- 5. Click **OK** to save the settings.

Configuring source data processing options

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



- 4. Specify how the source will process data.
 - Replication Packets to 1 Mirror Packet—You can specify the ratio of replication
 packets to mirror packets that are placed in the source queue. Specify a larger number if
 you have a busy network that has heavy replication. Also, if you anticipate increased
 network activity during a mirror, increase this number so that the replication queue does not
 get too large.
 - Replicate NT Security by Name—This is a Windows option only.
 - **Ignore Delete Operations**—This option allows you to keep files on the target machine after they are deleted on the source. When a file is deleted on the source, that delete operation is not sent to the target. (All edits to files on the source are still replicated to the

target; only deletions of whole files are ignored.) This option may be useful to give you an opportunity to make a backup of these files in the event they are needed in the future.



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

- Maximum Pending Mirror Operations—This option is the maximum number of mirror operations that are queued on the source. The default setting is 1000. If, during mirroring, the mirror queued statistic regularly shows low numbers, for example, less than 50, this value can be increased to allow Double-Take 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.

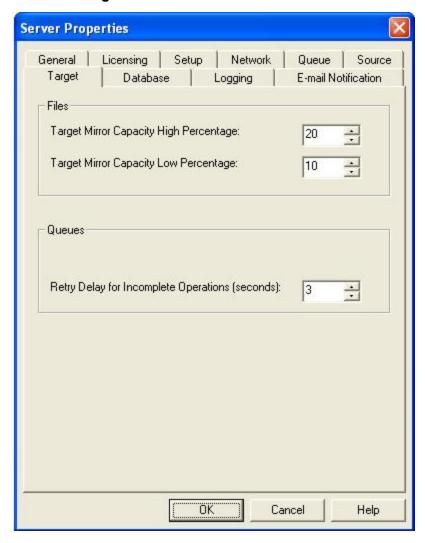


Database applications may update files without changing the date, time, or file size. Therefore, if you are using database applications, you should use the Block Checksum All option to ensure proper file comparisons.

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

Configuring target data processing options

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



- 4. Specify how the target will process data.
 - Target Mirror Capacity High Percentage—You can specify the maximum percentage of system memory that can contain mirror data before the target signals the source to pause the sending of mirror operations. The default setting is 20.
 - Target Mirror Capacity Low Percentage—You can specify the minimum percentage of system memory that can contain mirror data before the target signals the source to resume the sending of mirror operations. The default setting is 10.
 - Retry Delay for Incomplete Operations (seconds)—This option specifies the amount of time, in seconds, before retrying a failed operation on the target. The default setting is 3.
- 5. Click **OK** to save the settings.

Specifying the Double-Take database storage files

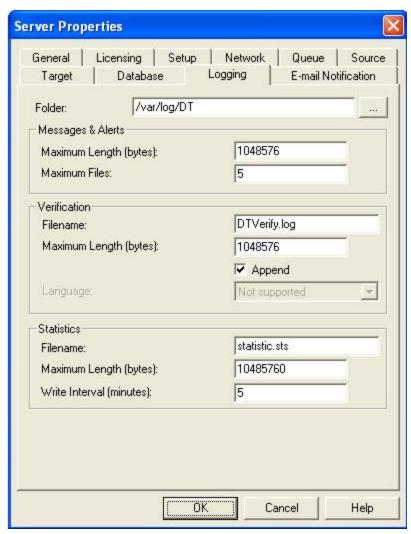
- 1. Right-click a server on the left pane of the Replication Console.
- 2. Select Properties
- Select the **Database** tab.



- 4. Specify the database files that store the Double-Take replication set, connection, and scheduling information.
 - **Folder**—Specify the directory where each of the database files on this tab are stored. The default location is the directory where the Double-Take program files are installed.
 - Replication Set—This database file maintains which replication sets have been created on the server along with their names, rules, and so on. The default file name is DblTake.db.
 - Connection—This database file maintains the active source/target connection information. The default file name is connect.sts.
 - **Schedule**—This database file maintains any scheduling and transmission limiting options. The default file name is schedule.sts.
- 5. Click **OK** to save the settings.

Specifying file names for logging and statistics

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



- 4. Specify the location and file names for the log and statistics files.
 - **Folder**—Specify the directory where each of the log files on this tab are stored. The default location is the directory where the Double-Take 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 name of the verification log, which is by default 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.

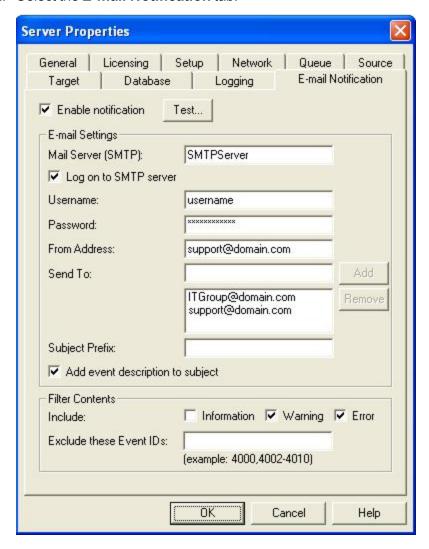
Statistics

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

E-mailing system messages

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

- 1. To enable e-mail notification for a server, right-click the server in the left pane of the Replication Console and select **Properties**.
- 2. Select the **E-mail Notification** tab.



3. Select Enable notification.



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

4. Specify your e-mail settings.

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



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

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

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

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

If desired, enable **Add event description** to subject to have the description of the message appended to the end of the subject line. This field is optional.

• Filter Contents—Specify which messages that you want to be sent via e-mail. Specify Information, Warning, and/or Error. You can also specify which messages to exclude based on the message ID. Enter the message IDs as a comma or semicolon separated list. You can indicate ranges within the list.



You can test e-mail notification by specifying the options on the E-mail Notification tab and clicking **Test**. If desired, you can send the test message to a different e-mail address by selecting **Send To** and entering a comma or semicolon separated list of addresses. Modify the message text up to 1024 characters, if necessary. Click **Send** to test the e-mail notification. The results will be displayed in a message box.



Click **OK to** close the message and click **Close** to return to the E-mail Notification tab

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

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

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

Chapter 15 Security

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. The groups can be local or LDAP (Lightweight Directory Access Protocol). To gain access to a particular Double-Take source or target, the user must provide a valid local user account that is 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. This security group name is dtamin, and the default group ID is 501.
- **Monitor Access**—Servers and statistics can be viewed, but functionality is not available. This security group name is dtmon, and the default group ID is 502.
- No Access—Servers appear in the clients, but no access to view the server details is available.

Although Double-Take passwords are encrypted when they are stored, Double-Takesecurity 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.

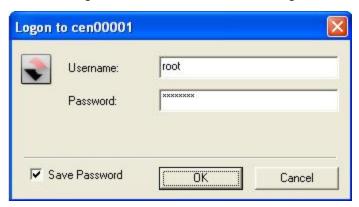
Logging on and off

To ensure protection of your data, Double-Take offer multi-level security using native operating system security features. Privileges are granted through membership in user groups defined on each machine running Double-Take. 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 features are available for that machine.
- Monitor Access—Servers and statistics can be viewed, but functionality is not available.
- No Access—Servers appear in the clients, but no access to view the server details is available.

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

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



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



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

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

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

Double-Take 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 Replication Console and disable **Use ICMP to verify server availability**.

Administrator rights 38

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

Monitor rights 🎘

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

No rights 🏥

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

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

Volume 2 Linux full server to ESX appliance protection

Create a full server to ESX appliance job when you want to protect an entire physical server or virtual machine to an ESX target. There is no reverse protection for this job. Linux full server to ESX appliance protection is distinct from Linux files and folders protection. For example, it has its own requirements and console. Use the following links to access information and steps specific to full server to ESX appliance protection.

- Getting started on page 213
- Managing and controlling full server to ESX appliance jobs on page 250
- Failing over full server to ESX appliance jobs on page 274
- Double-Take Console on page 275
- Managing servers on page 283

Chapter 1 Getting started

For full server to ESX appliance protection, you will need to complete the following steps, in order.

- 1. Review the *Full server to ESX appliance requirements* on page 214 to make sure your environment meets the requirements.
- 2. Install Double-Take on your virtual recovery appliance. See *Installing the Double-Take virtual recovery appliance* on page 219.
- 3. Install the Double-Take Console on a Windows machine. See *Installing the Double-Take Console* on page 223.
- 4. Add your virtual recovery appliance to the Double-Take Console. See *Adding a virtual recovery appliance to the console* on page 225.
- 5. License your virtual recovery appliance. See *Licensing the Double-Take virtual recovery appliance* on page 228.
- 6. Install Double-Take on your Linux source servers. You can install directly on your source server using the .rpm files located in the service/uvra/installers/linux/ directory on your appliance or the files from the DVD or web download (see *Installing Double-Take Availability on Linux* on page 229), or you can use the Double-Take Console to push the installation out to your source (see *Installing using the Double-Take Console* on page 231).
- Add your source servers to your Double-Take Console. See Adding a Linux source to the console on page 233.
- 8. Create your full server to ESX appliance job. See *Creating a full server to ESX appliance job* on page 236.

Once your job is created and running, see the following sections to manage your job.

- Managing and controlling full server to ESX appliance jobs on page 250—You can view status information about your job and learn how to control the job.
- Failing over full server to ESX appliance jobs on page 274—Use this section when a failover condition has been met or if you want to failover manually.

Full server to ESX appliance requirements

Use these requirements for full server to ESX appliance protection.

- Source server—The source server can be a physical or virtual server running any of the following operating systems.
 - Operating system—Red Hat Enterprise Linux, Oracle Enterprise Linux, and CentOS
 - Version—5.8 through 5.9
 - Kernel type for x86 (32-bit) architectures—Default, SMP, Xen, PAE
 - Kernel type for x86-64 (64-bit) architectures—Default, SMP, Xen
 - File system—Ext3, Ext4, XFS
 - Operating system——Red Hat Enterprise Linux, Oracle Enterprise Linux, and CentOS
 - Version—6.3 through 6.4
 - Kernel type for x86 (32-bit) architectures—Default
 - Kernel type for x86-64 (64-bit) architectures—Default
 - File system—Ext3, Ext4, XFS (64-bit only)
 - Operating system——Ubuntu
 - Version—10.04.3
 - Kernel version—2.6.32-33
 - Kernel type for x86 (32-bit) architectures—386, Generic, Generic-pae
 - Kernel type for x86-64 (64-bit) architectures—Generic, Preempt, Server
 - File system—Ext2, Ext3, Ext4, XFS
 - Operating system——Ubuntu
 - Version—10.04.4
 - Kernel version—2.6.32-38
 - Kernel type for x86 (32-bit) architectures—386, Generic, Generic-pae
 - Kernel type for x86-64 (64-bit) architectures—Generic, Preempt, Server
 - File system—Ext2, Ext3, Ext4, XFS
 - Operating system——Ubuntu
 - Version—12.04.0
 - Kernel version—3.2.0-23
 - Kernel type for x86 (32-bit) architectures—Generic, Virtual, Generic-pae, Lowlatency
 - Kernel type for x86-64 (64-bit) architectures—Generic, Virtual, Lowlatency
 - File system—Ext2, Ext3, Ext4, XFS
 - Operating system——Ubuntu
 - Version—12.04.1
 - Kernel version—3.2.0-29
 - Kernel type for x86 (32-bit) architectures—Generic, Generic-pae, Virtual

- Kernel type for x86-64 (64-bit) architectures—Generic, Virtual
- File system—Ext2, Ext3, Ext4, XFS
- Operating system——Ubuntu
 - Version—12.04.2
 - Kernel version—3.5.0-23
 - Kernel type for x86 (32-bit) architectures—Generic
 - Kernel type for x86-64 (64-bit) architectures—Generic
 - File system—Ext2, Ext3, Ext4, XFS
- Operating system——Ubuntu
 - **Version**—12.04.3
 - Kernel version—3.8.0-29
 - Kernel type for x86 (32-bit) architectures—Generic
 - Kernel type for x86-64 (64-bit) architectures—Generic
 - File system—Ext2, Ext3, Ext4, XFS



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

Oracle Enterprise Linux support is for the mainline kernel only, not the Unbreakable kernel.

Double-Take does not support stacking filesystems, like eCryptFS.

- Packages and services—Your Linux server must have the following packages and services
 installed before you can install and use Double-Take Availability. See your operating system
 documentation for details on these packages and utilities.
 - sshd (or the package that installs sshd)
 - Isb
 - parted
 - /usr/bin/which
 - /usr/sbin/dmidecode
 - /usr/bin/scp (only if you will be performing push installations from the Double-Take Console to your source servers)
- **Target host server**—The target host server must be an ESX server. It can be any of the following ESX operating systems.
 - ESX 4.1 Standard, Advanced, Enterprise, or Enterprise Plus
 - ESXi 4.1 Standard, Advanced, Enterprise, or Enterprise Plus
 - ESXi 5.0 Standard, Enterprise, or Enterprise Plus
 - ESXi 5.1 Standard, Enterprise, or Enterprise Plus
 - ESXi 5.5 Standard, Enterprise, or Enterprise Plus

- Virtual recovery appliance—The target ESX host must have an existing virtual machine, known as a virtual recovery appliance. This will be an OVF (Open Virtualization Format) virtual machine included with Double-Take. You must install this virtual machine before you can establish protection. When you establish protection, the virtual recovery appliance will create a new virtual machine, mount disks, format disks, and so on. If failover occurs, the new virtual machine is detached from the virtual recovery appliance and powered on. Once the new virtual machine is online, it will have the identity, data, and system state of the source. Since the virtual recovery appliance maintains its own identity, it can be reused for additional failovers. Keep in mind the following caveats for the virtual recovery appliance.
 - The virtual recovery appliance must be a standalone virtual machine.
 - It should not reside in any multiple virtual machine vApp.
 - The appliance is pre-configured for optimal performance. You do not need to modify the memory, CPU, or other configurations.
 - You should not install or run anything else on this appliance.
 - The firewall is disabled and should remain disabled.
 - A single virtual recovery appliance can protect a maximum of 59 volume groups and raw block devices (combined) from any number of sources.
- vCenter—vCenter is not required, but if you are using it, then you must use version 4.1 or later.
- vCenter permissions—If you want to limit the permissions required for the vCenter account that you will be using for your full server to ESX appliance job, your account must have at a minimum the permissions listed below.
 - Datastore—Allocate Space, Browse Datastore, Low level file operations, and Remove File
 - Host, Local Operations—Create Virtual Machine, Delete Virtual Machine, and Reconfigure virtual machine
 - Network—Assign Network
 - Resource—Assign virtual machine to resource pool
 - Scheduled Task—Create Tasks, Modify Task, Remove Task, and Run Task
 - Tasks—Create task and Update task
 - Virtual Machine, Configuration—Add existing disk, Add new disk, Add or remove device, Change resource, Modify device settings, and Remove disk
 - Virtual Machine, Interaction—Device connection, Power off, and Power on
 - Virtual Machine, Inventory—Create new, Register, Remove, and Unregister
- vMotion—Host vMotion is only supported if you are using vCenter. Storage vMotion is not supported.
- **System memory**—The minimum system memory on each server should be 1 GB. The recommended amount for each server is 2 GB.
- Disk space for program files—This is the amount of disk space needed for the Double-Take program files. This is approximately 285 MB on a Linux source server. The appliance needs approximately 620 MB.



Approximately 45 MB will be on your /root partition, 40 MB on your /usr partition, and 200 MB in /opt.

Make sure you have additional disk space for Double-Take queuing, logging, and so on.

- **Server name**—Double-Take includes Unicode file system support, but your server name must still be in ASCII format. If you have the need to use a server's fully-qualified domain name, your server cannot start with a numeric character because that will be interpreted as an IP address. Additionally, all Double-Take servers and appliances must have a unique server name.
- Protocols and networking—Your servers must meet the following protocol and networking requirements.
 - Your servers must have TCP/IP with static IP addressing.
 - IPv4 is the only supported version.
- Name resolution—Your servers must have name resolution or DNS. The Double-Take Console must be able to resolve the virtual recovery appliance, and the virtual recovery appliance must be able to resolve all source servers. For details on name resolution options, see your Linux documentation or online Linux resources.
- **Ports**—Port 1501 is used for localhost communication. Ports 1500, 1505, 1506, 6325, and 6326 are used for component communication and must be opened on any firewall that might be in use.
- **Security**—Double-Take security is granted through membership in user groups. The groups can be local or LDAP (Lightweight Directory Access Protocol). A user must provide a valid local account that is a member of the Double-Take security groups
- **SELinux policy**—SELinux should be disabled on the source.
- Snapshots—Double-Take snapshots are not supported with full server to ESX appliance jobs.
- **Supported configurations**—The following table identifies the supported configurations for a full server to ESX appliance job.

Configuration		Supported	Not Supported
Source to target configuration ¹	One to one, active/standby	Х	
	One to one, active/active		Х
	Many to one	Х	
	One to many		Х
	Chained		Х
	Single server		Х

Configuration		Supported	Not Supported
Server configuration	Standalone to standalone	Х	
	Standalone to cluster		Х
	Cluster to standalone		Х
	Cluster to cluster		Х
Upgrade configuration ²	Upgrade 4.7 Replication Console connection to 7.0 full server to ESX appliance job		Х
	Upgrade 6.0 full server to ESX appliance job to 7.0 full server to ESX appliance job		Х
	Version 7.0 console can view or create job for 4.7 source and 4.7 target		Х
Version 7.0 console	Version 7.0 console can create job for 6.0 source and 6.0 target		Х
	Version 7.0 console can create job for 7.0 source and 7.0 target	Х	

- 1. See *Supported configurations* on page 11 for details on each of the source to target configurations.
- 2. When upgrading from version 4.7, you will have to delete the existing connection before the upgrade and create a new job after the upgrade. When upgrading from version 6.0, delete the existing 6.0 job, upgrade your source, deploy a new appliance, and then create a new 7.0 job.

Installing the Double-Take virtual recovery appliance

If you will be creating full server to ESX appliance jobs, you will need to install the Double-Take virtual recovery appliance. You have the choice of installing an OVF (Open Virtualization Format) prepackaged virtual machine or installing the software on your own pre-formatted virtual machine.



If you already have an existing Double-Take virtual recovery appliance that you want to upgrade, use the instructions *Installing the appliance manually* on page 221 to install the latest .deb files on your existing appliance.

Deploying the OVF virtual machine

- 1. Make sure you review the *Full server to ESX appliance requirements* on page 214 before you begin the installation.
- 3. From your vSphere client, select the root of your tree, select **File**, **Deploy OVF Template**, and specify the dtvra_<version_number>.ova file. Click **Next** to continue.
- 4. Review the OVF template details and click **Next** to continue.
- 5. If you are connected to a vCenter, review the Vision Solutions license agreement. You must accept the license agreement by clicking **Accept**. If you do not have vCenter, you will need to accept the license agreement when the appliance first starts. Click **Next** to continue.
- 6. Specify the name and location where you want to install the appliance. This should be the same datacenter where your target ESX host is located. Click **Next** to continue.
- 7. Select the ESX host where you want to install the appliance. This is your target ESX host. Click **Next** to continue.
- 8. If desired, select the resource pool where you want to install the appliance. You will not have this option if you do not have resource pools. Click **Next** to continue.
- 9. Select the datastore where you want to store the appliance files. Make sure the datastore has at least 16 GB of space available. Click **Next** to continue.
- Select the type of disk to create. The names you see displayed will be dependent on your vSphere client.
 - Flat or Thick Provision Lazy Zeroed—This disk type allocates the full amount of the disk space immediately, but does not initialize the disk space to zero until it is needed. This option is only available for ESX version 5 and later.
 - Thick or Thick Provision Eager Zeroed—This disk type allocates the full amount of the disk space immediately, initializing all of the allocated disk space to zero.
 - Thin Provision—This disk type does not allocate the disk space until it is needed.
- 11. Click **Next** to continue.
- 12. If your selected ESX host has multiple networks, map the network on the appliance to a network on your ESX host, and click **Next** to continue.

- 13. If you are connected to a vCenter, assign the virtual machine and networking properties. If you do not have vCenter, you will need to configure these settings after the appliance starts.
 - Root password—Set a password for the root login. If you do not make a change, the default password will be used. The default password is vision.
 - **Time zone setting**—Select the virtual machine's time zone. The default setting is UTC.
 - **Hostname**—Specify the hostname or full qualified domain name to assign to the virtual machine. If you leave this option blank, Double-Take will use reverse DNS lookup for the identification.
 - **Default gateway**—Specify the virtual machine's default gateway.
 - **DNS server**—Specify the virtual machine's DNS server.
 - **Network IP address**—Specify the IP address to assign to the virtual machine. You must use an IPv4 address.
 - Network subnet mask—Specify the subnet mask for the assigned IP address.



vCenter 4.x will display your deployment password in visible text. To avoid a security risk, you should use the default password or set a temporary password, and then after the appliance deployment is complete, you should change the default or temporary password to a new secure password.

- 14. Click Next to continue.
- 15. Review the **Deployment settings**. If you want to make any changes, click **Back**. If not, enable the **Power on after deployment** option (which is only available for ESX version 5 and later), and click **Finish**.
- 16. After the appliance deployment is complete, click **Close**.

After the appliance has been installed, you will need to install the Double-Take Console on a Windows machine, if you haven't already. You can use the .exe file located in the service/uvra/installers/windows/directory on your appliance, or see the instructions *Installing the Double-Take Console* on page 223 to install the console from the DVD or a web downloaded file. In either case, make sure you select the **Client Components Only** installation option. After you have installed the Double-Take Console, continue with adding your virtual recovery appliance to the console on page 225.



Once the appliance has been installed and configured, you should not need to access or use it. However, if you need to change any network settings, you can modify them using the **Configure Network** menu on the appliance or by using the web interface by going to https://ip_address:5480, where ip_address is the current IP address of the appliance. Keep in mind if you change the network settings using the web interface, you will lose your connection to that interface because it will still be using the original IP address.

The appliance menu runs as root, so make sure you take advantage of VMware security to restrict access to the appliance. See your VMware documentation for details on limiting access to the appliance's console.

Installing the appliance manually

Before you begin the manual appliance installation, you must have an existing VMware hosted virtual machine that meets the following requirements.

- The virtual machine must be running Ubuntu 10.04.x AMD64 edition.
- The virtual machine must have at least 4 GB of virtualized physical RAM.
- The virtual machine must have at least two CPUS (two virtual sockets, not two virtual cores).
- The virtual machine must have at least 16 GB of disk space available.
- The virtual machine must have a valid, working network configuration, including DNS.
- If your source servers that you will be protecting have the XFS file system, the virtual machine
 must have the xfsprogs and xfsdump packages installed.
- The virtual machine must be dedicated to Double-Take processing only. Do not use the virtual machine for any other activity (web server, database server, and so on).
- Make sure you review the Full server to ESX appliance requirements on page 214 for any other requirements not specifically defined here.
- Locate the .deb installation file either on your Double-Take DVD or from the location where you
 downloaded it. For example, for version 7.0, the installation file might be doubletake_
 7.0.0.1124.0_amd64.deb If necessary, copy the installation file to a location that is accessible
 from your virtual machine.
- 2. Install the .deb file on your virtual machine.
- 3. If you are going to push the Double-Take installation from your appliance to your source servers, you will also need to install the installers package. Locate the .deb installers file either on your Double-Take DVD or from the location where you downloaded it. For example, for version 7.0, the installation file name might be doubletake-installers_7.0.0.1225.0_all.deb. Note that the version number will be different than the first .deb file. If necessary, copy this file to a location that is accessible from your virtual machine.
- 4. Install the installers .deb file on your virtual machine.
- 5. If you are going to push the Double-Take installation from your appliance to your source servers, the root user must also use the SSH client at least once before push installations can be completed. You can establish an ssh connection from the appliance to any other server or to itself. This only has to be done once on the appliance.
- 6. After the installation your virtual machine is complete, you will need to accept the license agreement and start the Double-Take service.
 - Run the DTSetup command from the shell prompt to start DTSetup. The command is casesensitive.



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.

b. Review the Vision Solutions license agreement and accept the terms of agreement by typing yes.

- c. When the DTSetup menu appears, enter the number of the menu option you want to access. In this case, select **Start/Stop Double-Take daemon**.
- d. Select **Start Double-Take and process driver config** to start the DT daemon and load the Double-Take drivers.
- e. After Double-Take is running, press Q as many times as needed to return back to the main menu and then exit DTSetup.

After the appliance has been installed, you will need to install the Double-Take Console on a Windows machine, if you haven't already. You can use the .exe file located in the service/uvra/installers/windows/directory on your appliance (if you installed the installers package), or see the instructions *Installing the Double-Take Console* on page 223 to install the console from the DVD or a web downloaded file. In either case, make sure you select the **Client Components Only** installation option. After you have installed the Double-Take Console, continue with adding your virtual recovery appliance to the console. See *Adding a virtual recovery appliance to the console* on page 225.

Installing the Double-Take Console

The Double-Take Console must be installed on a Windows machine. See *Double-Take Console requirements* on page 276 for specific details.

- 1. Close any open applications.
- 2. Start the installation program using the appropriate instructions, depending on your media source.
 - Physical media—If auto-run is enabled, the installation program will start automatically.
 To manually start the program, run autorun.exe from your physical media.
 - **Web download**—Launch the .exe file that you downloaded from the web.
- When the installation program begins, the Autorun appears allowing you to install software and view documentation and online resources. To install the console, , select the Install Double-Take Availability link.
- 4. Depending on your version of Windows and the components you have installed, you may see an initial screen indicating that you need to install or enable Microsoft .NET Framework. If you do not see this screen, your server already has the appropriate version of Microsoft .NET. You must install or enable Microsoft .NET before installing Double-Take. Select **Yes** to install or enable Microsoft .NET and click **Continue**.
- 5. You will be given the opportunity to check for a more recent version of the software.
 - If you do not want to check for a later version, select No and click Next.
 - If you want to check for a later version, select **Yes** and click **Next**. The installation program will establish an Internet connection from your server to the Vision Solutions web site.
 - If later versions are found, they will be listed. Highlight the version you want and
 either download that version and install it automatically or download that version and
 exit the installation. (If you exit the installation, you can run the updated installation
 later directly from the location where you saved it.)
 - If no later versions are found, continue with the current installation.
 - If an Internet connection cannot be established, continue with the current installation or install a previously downloaded version.
- 6. Review the Vision Solutions license agreement. You must scroll through and review the entire license agreement. You must accept the license agreement in order to continue with the installation program. Click **Next** to continue.
- Review the activation notice. Most Double-Take licenses require activation after installation for full
 product functionality. Failure to activate licenses that require it will cause your Double-Take jobs to
 fail. See *Licensing a server or appliance* on page 298 for more details.
- 8. Click **OK** to continue.
- 9. Select the **Client Components Only** installation option. The client components do not require an activation code, but are required to administer Double-Take servers throughout the organization.
- If desired, specify where the Double-Take files will be installed by clicking Change, specifying a location, and then clicking OK.
- 11. Click Next to continue.
- 12. If the machine where you are installing has Windows Firewall enabled, you will be given the opportunity to open and reassign any firewall ports for Double-Take use.
 - Open only the ports that are not in use—This option will open any firewall ports that are not in use. The ports that are opened will be assigned to Double-Take.

- Open all ports, reassigning the ports in use to Double-Take—This option will open all necessary firewall ports, reassigning any to Double-Take as needed.
- Do not configure the Windows Firewall at this time—This option will not modify any firewall ports. If you select this option, you will have manually modify your firewall settings for Double-Take processing.
- 13. Click **Next** to continue.
- 14. If you are satisfied with the selections you have made and are ready to begin copying the Double-Take files, click **Install**.
- 15. After the files have completed copying, click **Finish** to exit the installation program.

Adding a virtual recovery appliance to the console

You must insert your virtual recovery appliance into your console session before inserting your source servers because a source server needs to be associated with an appliance that is already inserted in the console.

- 1. Select **Get Started** from the toolbar.
- 2. Select Add servers and click Next.
- 3. On the **Manual Entry** tab, specify the appliance information.
 - Server—This is the name or IP address of the appliance to be added to the console. See the following NAT configuration section if you have a NAT environment.
 - **User name**—Specify the root user or another user that has the administrator role on the appliance.
 - Password—Specify the password associated with the User name you entered.
- 4. You can expand the **More Options** section to configure the following settings.
 - Domain—If you are working in a domain environment, specify the Domain.
 - Associate with Double-Take Linux Appliance—Leave this field blank. You will be specifying this field when you insert your Linux source sever.
- 5. After you have specified the appliance information, click **Add**.
- 6. Repeat steps 3 through 5 for any other appliances you want to add.
- 7. If you need to remove appliances from the list of **Servers to be added**, highlight it and click **Remove**. You can also remove all of them with the **Remove All** button.
- 8. When your list of **Servers to be added** is complete, click **OK**.

NAT configuration

The name or IP address you use to add a server or appliance to the console is dependent on where are you are running the console. For a full server to ESX appliance job, the console only communicates with the appliance. Therefore, you need to use the address which the appliance has access to when adding the source server. Use the following table to determine what name or IP address to enter depending on the location where you are running the console.



In this table, public addresses are those addresses that are publicly available when a server is behind a NAT router. Private addresses are those addresses that are privately available when a server is behind a NAT router. An address that is not labeled as public or private are for servers that are not behind a NAT router. This is generally a public address but is not named as such in this table to try to more clearly identify when a public NAT address needs to be used.

Location of source and appliance	Location of Double-Take Console	How to add the source and appliance to the Double-Take Console	
If your source and appliance are behind individual NAT routers,	and your Double-Take Console is located behind the NAT router with the source,	specify the public IP address of the source (which is the public IP address of the source's NAT router) and the public IP address of the appliance (which is the public IP address of the appliance's NAT router).	
	and your Double-Take Console is located behind the NAT router with the appliance,	specify the public IP address of the source (which is the public IP address of the source's NAT router) and the name or private IP address of the appliance.	
	and your Double-Take Console is located between the two NAT routers,	specify the public IP address of the source (which is the public IP address of the source's NAT router) and the public IP address of the appliance (which is the public IP address of the appliance's NAT router).	
	and your Double-Take Console is located behind a third NAT router,		

If your source is behind a NAT router but your appliance is not,	and your Double-Take Console is located behind the NAT router with the source,	specify the public IP address of the source (which is the public IP address of the source's NAT router) and the name or IP address of the appliance.
	and your Double-Take Console is located on the appliance network,	specify the public IP address of the source (which is the public IP address of the source's NAT router) and the name or IP address of the appliance.
If your appliance is behind a NAT router but your source is not,	and your Double-Take Console is located behind the NAT router with the appliance,	specify the name or IP address of the source and the name or private IP address of the appliance.
	and your Double-Take Console is located on the source network,	specify the name or IP address of the source and the public IP address of the appliance (which is the public address of the appliance's NAT router).
If your source and appliance are both behind a single NAT router,	and your Double-Take Console is located outside of the router,	specify the name or private IP address of the source and the public IP address of the appliance (which is the public address of the apliance's NAT router).



As noted above, make sure you insert your virtual recovery appliance into your console session before inserting your source servers because a source server needs to be associated with an appliance that is already inserted in the console.

Licensing the Double-Take virtual recovery appliance

- 1. Make sure you have added your appliance to the console. See *Adding a virtual recovery appliance to the console* on page 225.
- 2. Highlight your appliance on the **Manage Servers** page and click **View Server Details** from the toolbar. Under **Tasks**, select **Edit server properties**.
- 3. In the box under Add activation codes and activation keys, enter your activation code and click Add. Some activation codes will have a 14-day grace period to activate the license. If your license requires activation as noted on the Edit Server Properties page, see *License management and activation* on page 295 for details on license activation.
- 4. Click **OK** to return to the **Manage Servers** page.

Once your appliance is licensed, you will need to install Double-Take on your source servers. You can use the .rpm files located in the service/uvra/installers/linux/ directory on your appliance or the files from the DVD or web download. See *Installing Double-Take Availability on Linux* on page 229 for installation instructions. Alternatively, you can use the Double-Take Console to push the installation out to your source servers. See *Installing using the Double-Take Console* on page 231.

Installing Double-Take Availability on Linux

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



Make sure you review the *Full server to ESX appliance requirements* on page 214 before you begin the installation.

The instructions below are formatted for Red Hat and similar operating systems. You may need to modify the commands to fit your operating system. For example, on Ubuntu you should use dpkg -i to install the .deb file. See your operating system reference documentation for appropriate installation commands for your server.

- 1. Locate the installation files. The .rpm files are located in the service/uvra/installers/linux/ directory on your appliance, on the DVD, or in the web download.
- 2. Determine the installation package that is appropriate for your operating system. There are separate .rpm installation packages for 32-bit and 64-bit architectures. Make sure you are installing the correct .rpm file. If you are uncertain about the architecture of your machine, you can use the uname -m command to determine it. Additionally, the version_numbers in the installation file name will vary and will correspond to the version of Double-Take Availability you are installing.
 - 32-bit architecture—DoubleTake-version_numbers.i386.rpm
 - 64-bit architecture—DoubleTake-version_numbers.x86_64.rpm

For example, if you are installing version 7.0.0.0.1124, the installation files would be DoubleTake-7.0.0.0-1124.i386.rpm or DoubleTake-7.0.0.0-1124.x86 64.rpm.

- 3. If you are upgrading an existing Double-Take Availability server, you should complete the following steps before upgrading.
 - Shutdown the protected application(s)
 - Stop DT, the Double-Take daemon
 - Detach DTLOOP, if it is being used
- 4. Once you have determined the appropriate installation package to use, you can install the software from the user interface or from the command line.
 - **UI installation**—Double-click the .rpm file from the UI and confirm the installation or upgrade.
 - Command line installation—Use the following steps to install from a command line.
 - a. Make sure you are a root/uid 0 user.
 - b. Go to a shell prompt by launching a terminal application from your UI or logging in via the Linux virtual console.
 - c. If you are installing from a DVD, 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-7.0.0.0-1124.i386.rpm.

A successful installation returns you to the shell prompt.

- 5. After the installation on the Linux server is complete, you will need to license your server using the following instructions.
 - a. Run the DTSetup command from the shell prompt to start DTSetup. The command is case-sensitive.



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

- b. If prompted, review the Vision Solutions license agreement and accept the terms of agreement by typing yes.
- c. When the DTSetup menu appears, enter the number of the menu option you want to access. In this case, select **Setup tasks**, which will allow you to activate your server.
- d. Select Set Activation Code Menu.
- e. Select Set Activation Code in /etc/DT/DT.conf.
- f. Enter your activation code and press Enter. The activation code will automatically be inserted into the configuration file.
- g. Press Enter to return to the menu.
- h. Press Q as many times as needed to return back to the main menu and then exit DTSetup.

Installing using the Double-Take Console

You can use the Double-Take Console to install or upgrade Double-Take Availability on your source servers. The installation is a full server installation.



The push installation process uses a combination of SSH (secure shell) and SCP (secure copy) to copy the installation binary files and execute them. SCP is not installed by default on certain configurations of version 6.x of Red Hat Enterprise Linux, Oracle Enterprise Linux, and CentOS. You will need to install it manually to use the push installation process. See your Linux documentation for installation details. See *Full server to ESX appliance requirements* on page 214 for complete requirement details.

You can only push the installation process. You cannot push an upgrade.

- 1. Add the servers where you want to install Double-Take to your console session. See *Adding a Linux source to the console* on page 233.
- 2. From the **Manage Servers** page, highlight all of the servers where you want to install Double-Take, and select **Install** from the toolbar.
- 3. Each server needs an activation code for the installation.
 - a. If you have a single activation code that can be used on multiple servers, such as a site license or an evaluation code, highlight all of the servers where you want to apply the same activation code and click **Set Activation codes**. If you have unique activation codes for each of your servers, highlight just one server and click **Set Activation codes**.
 - b. Type in your activation code or click **Choose from inventory** to open the Activation Codes dialog box where you can select the activation code you want to apply. For more information on the license inventory, see *Managing the Double-Take license inventory* on page 296.
 - c. Click **OK** to configure the installation or upgrade for the selected activation code.
 - d. Repeat steps a-c for any additional servers that are using unique activation codes.



If you are pushing out a code that requires activation, you will have 14 days to activate the code. See *License management and activation* on page 295 for more details on license activation.

4. The **Default Installation Options** section contains the default settings from the console's **Options** page. These settings will be applied to all of the servers you are installing to or upgrading. If desired, modify any of the installation options. See *Console options* on page 277.



Options you modify when using the push installation will be replaced by the default values from the **Options** page each time you use the push installation.

 Location of 32-bit installation package—This option does not apply to Linux source servers.

- Location of 64-bit installation package—This option does not apply to Linux source servers.
- Location of .NET 3.5 SP1 installation package—This option does not apply to Linux source servers.
- Temporary folder for installation package—Specify a temporary location (on the server where you are installing Double-Take) where the installation files will be copied and run. You need approximately 130 MB of space in the specified location.
- Installation folder—Specify the location where you want to install Double-Take on each server.
- Queue folder—Specify the location where you want to store the Double-Take disk queue on each server.
- Amount of system memory to use—Specify the maximum amount of memory, in MB, that can be used for Double-Take processing.
- Minimum free disk space—This is the minimum amount of disk space in the specified
 Queue folder that must be available at all times. This amount should be less than the
 amount of physical disk space minus the disk size specified for Limit disk space for
 queue.
- **Do not use disk queue**—This option will disable disk queuing. When system memory has been exhausted, Double-Take will automatically begin the auto-disconnect process.
- Unlimited disk queue—Double-Take will use an unlimited amount of disk space in the specified Queue folder for disk queuing, which will allow the queue usage to automatically expand whenever the available disk space expands. When the available disk space has been used, Double-Take will automatically begin the auto-disconnect process.
- Limit disk space for queue—This option will allow you to specify a fixed amount of disk space, in MB, in the specified Queue folder that can be used for Double-Take disk queuing. When the disk space limit is reached, Double-Take will automatically begin the auto-disconnect process.
- 5. After you have configured your installation options, click **Install**.

Adding a Linux source to the console

You must have your virtual recovery appliance already inserted into your console session before inserting your source servers because a source server needs to be associated with an appliance. See *Adding a virtual recovery appliance to the console* on page 225 if you have not inserted your appliance.

- 1. Select **Get Started** from the toolbar.
- 2. Select Add servers and click Next.
- 3. On the **Manual Entry** tab, specify the source information.
 - **Server**—This is the name or IP address of the source server to be added to the console. See the following NAT configuration section if you have a NAT environment.
 - **User name**—Specify a user that is a member of the **DTAdmin** or **DTMon** security group on the server.
 - Password—Specify the password associated with the User name you entered.
- 4. You can expand the **More Options** section to configure the following settings.
 - **Domain**—If you are working in a domain environment, specify the **Domain**.
 - Associate with Double-Take Linux Appliance—Specify the name or IP address of your virtual recovery appliance. Your Linux source is now considered a proxied server.
- 5. After you have specified the server information, click **Add**.
- 6. Repeat steps 3 through 5 for any other servers you want to add.
- 7. If you need to remove servers from the list of **Servers to be added**, highlight it and click **Remove**. You can also remove all of them with the **Remove All** button.
- 8. When your list of **Servers to be added** is complete, click **OK**.

NAT configuration

The name or IP address you use to add a server or appliance to the console is dependent on where are you are running the console. For a full server to ESX appliance job, the console only communicates with the appliance. Therefore, you need to use the address which the appliance has access to when adding the source server. Use the following table to determine what name or IP address to enter depending on the location where you are running the console.



In this table, public addresses are those addresses that are publicly available when a server is behind a NAT router. Private addresses are those addresses that are privately available when a server is behind a NAT router. An address that is not labeled as public or private are for servers that are not behind a NAT router. This is generally a public address but is not named as such in this table to try to more clearly identify when a public NAT address needs to be used.

Location of source and appliance	Location of Double-Take Console	How to add the source and appliance to the Double-Take Console	
If your source and appliance are behind individual NAT routers,	and your Double-Take Console is located behind the NAT router with the source,	specify the public IP address of the source (which is the public IP address of the source's NAT router) and the public IP address of the appliance (which is the public IP address of the appliance's NAT router).	
	and your Double-Take Console is located behind the NAT router with the appliance,	specify the public IP address of the source (which is the public IP address of the source's NAT router) and the name or private IP address of the appliance.	
	and your Double-Take Console is located between the two NAT routers,	specify the public IP address of the source (which is the public IP address of the source's NAT router) and the public IP address of the appliance (which is the public IP address of the appliance's NAT router).	
	and your Double-Take Console is located behind a third NAT router,		

If your source is behind a NAT router but your appliance is not,	and your Double-Take Console is located behind the NAT router with the source,	specify the public IP address of the source (which is the public IP address of the source's NAT router) and the name or IP address of the appliance.
	and your Double-Take Console is located on the appliance network,	specify the public IP address of the source (which is the public IP address of the source's NAT router) and the name or IP address of the appliance.
If your appliance is behind a NAT router but your source is not,	and your Double-Take Console is located behind the NAT router with the appliance,	specify the name or IP address of the source and the name or private IP address of the appliance.
	and your Double-Take Console is located on the source network,	specify the name or IP address of the source and the public IP address of the appliance (which is the public address of the appliance's NAT router).
If your source and appliance are both behind a single NAT router,	and your Double-Take Console is located outside of the router,	specify the name or private IP address of the source and the public IP address of the appliance (which is the public address of the apliance's NAT router).

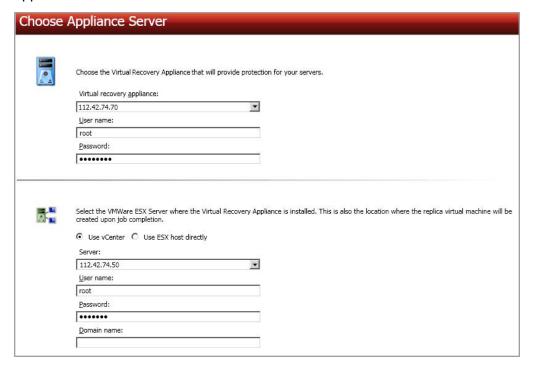


As noted above, make sure you insert your virtual recovery appliance into your console session before inserting your source servers because a source server needs to be associated with an appliance that is already inserted in the console.

Creating a full server to ESX appliance job

Use these instructions to create a full server to ESX appliance job.

- 1. Click Get Started from the toolbar.
- 2. Select Double-Take Availability and click Next.
- 3. Select Protect a Linux server with a Linux Virtual Recovery Appliance and click Next.
- 4. Choose your Double-Take virtual recovery appliance and the target ESX server that is hosting the appliance.



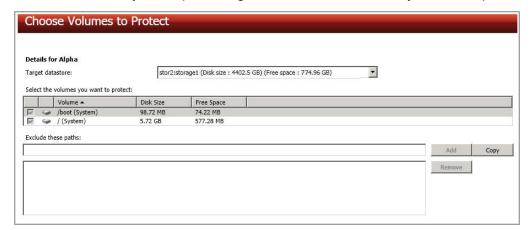
- Virtual Recovery Appliance—Specify the IP address of your appliance. You can also
 use the name of the appliance, if the Double-Take Console machine can resolve the
 appliance's IP address.
- User name—Specify the root user.
- **Password**—Specify the password that you assigned when first you powered on the appliance.
- Use vCenter—Select this option if you are using vCenter server.
- Use ESX host directly—Select this option if you are not using vCenter.
- Server—Specify the name or IP address of your vCenter server or ESX server.
- **User name**—If you are using vCenter, specify a user that has the administrator role for your vCenter. If you are using your ESX host directly, enter your ESX server credentials.
- Password—Specify the password associated with the User name you entered.
- 5. Click **Next** to continue.

6. Choose your source server. This is the Linux server that you want to protect.



- Current Servers—This list contains the servers currently available in your console session. Servers that are not licensed for the workflow you have selected will be filtered out of the list. Select your source server from the list.
- 7. Click **Next** to continue.

8. For the server that you are protecting, choose the volumes that you want to protect.



- Target datastore
 —Select a datastore on your ESX host where you want to store the
 replica virtual machine configuration files. You can specify the location of the .vmdk files
 later in the workflow.
- Select the volumes you want to protect—By default, the system and boot volumes will be selected for protection. You will be unable to deselect these volumes. Select any other volumes on the source that you want to protect.



The swap partition is excluded by default and you cannot select it, however, it will be created on the replica.

- Exclude these paths—If there are specific paths on a volume that you do not want to protect, specify those locations. Keep in mind that excluded data may impact the integrity of your applications. Specify any paths to exclude from protection and click Add. You can also specify a comma or semi-colon delimited list of paths and click Add, and separate exclude rules will be made for each path in your list. If you need to remove a path from being excluded, select the path and click Remove.
- 9. Click Next to continue.

10. You have many options available for your full server to ESX appliance job. Configure those options that are applicable to your environment.

Go to each page identified below to see the options available for that section of the **Set Options** page. After you have configured your options, continue with the next step on page 249.

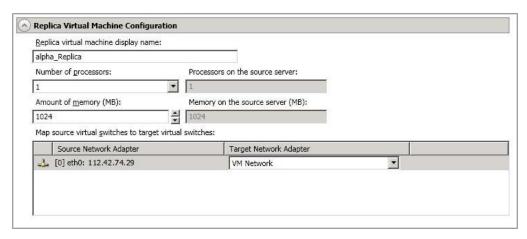
- General on page 240
- Replica Virtual Machine Configuration on page 241
- Replica Virtual Machine Volumes on page 242
- Replica Virtual Machine Network Settings on page 245
- Mirror, Verify & Orphaned Files on page 246
- Network Route on page 247
- Compression on page 248
- Bandwidth on page 249

General



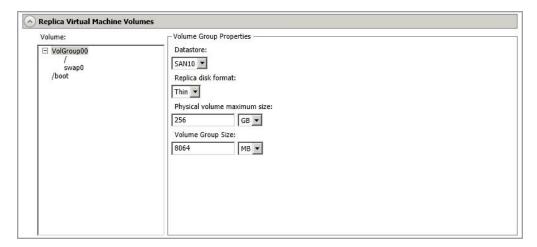
For the **Job name**, specify a unique name for your job.

Replica Virtual Machine Configuration



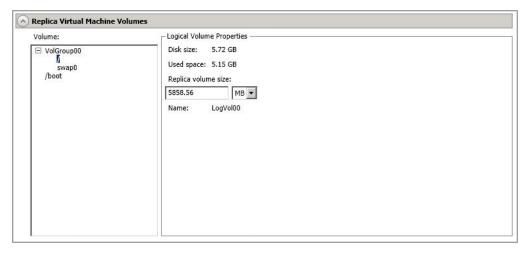
- **Replica virtual machine display name**—Specify the name of the replica virtual machine. This will be the display name of the virtual machine on the host system.
- Number of processors—Specify how many processors to create on the new virtual
 machine. The number of processors on the source is displayed to guide you in making an
 appropriate selection. If you select fewer processors than the source, your clients may be
 impacted by slower responses.
- Amount of memory—Specify the amount of memory, in MB, to create on the new virtual
 machine. The memory on the source is displayed to guide you in making an appropriate
 selection. If you select less memory than the source, your clients may be impacted by
 slower responses.
- Map source virtual switches to target virtual switches—Identify how you want to
 handle the network mapping after failover. The Source Network Adapter column lists the
 NICs from the source. Map each one to a Target Network Adapter, which is a virtual
 network on the target.

Replica Virtual Machine Volumes



If your source has volume groups, you will see them listed in the **Volume** list. Highlight a volume group and set the available **Volume Group Properties** that are displayed to the right of the **Volume** list.

- Datastore—Specify the datastore where you want to store the .vmdk files for the volume group. You can specify the location of the virtual machine configuration files on the previous Choose Volumes to Protect page.
- Replica disk format—Specify the format of the disk that will be created.
 - Flat—This disk format allocates the full amount of the disk space immediately, but
 does not initialize the disk space to zero until it is needed. This disk format is only
 available on ESX 5; if you select this disk type on ESX 4, a thick disk will be created.
 - Thick—This disk format allocates the full amount of the disk space immediately, initializing all of the allocated disk space to zero.
 - Thin—This disk format does not allocate the disk space until it is needed.
- Physical volume maximum size—Specify the maximum size, in MB or GB, of the virtual
 disks used to create the volume group. The default value is equal to the maximum size that
 can be attached to the datastore you selected. That will depend on your ESX version, your
 file system version, and the block size of your datastore.
- Volume Group size—Specify the maximum size, in MB or GB, of the volume group. The
 default value will match the source. This value cannot be less than the logical volumes total
 size that you are trying to create on the volume group.



If your source has logical volumes, you will see them listed in the **Volume** list. Highlight a logical volume and set the available **Logical Volume Properties** that are displayed to the right of the **Volume** list.



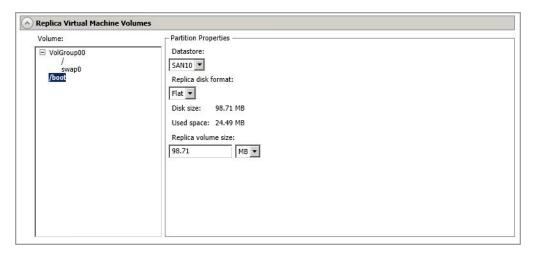
The size and space displayed may not match the output of the Linux df command. This is because df shows the size of the mounted file system not the underlying partition which may be larger. Additionally, Double-Take uses powers of 1024 when computing GB, MB, and so on. The df command typically uses powers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

- **Disk size**—This field displays the size of the logical volume on the source.
- **Used space**—This field displays the amount of disk space in use on the source logical volume.
- **Replica volume size**—Specify the size, in MB or GB, of the replica logical volume on the target. The value must be at least the size of the specified **Used space** on that volume.



In some cases, the replica virtual machine may use more virtual disk space than the size of the source volume due to differences in how the virtual disk's block size is formatted and how hard links are handled. To avoid this issue, specify the size of your replica to be at least 5 GB larger.

• Name—This field displays the logical volume name.



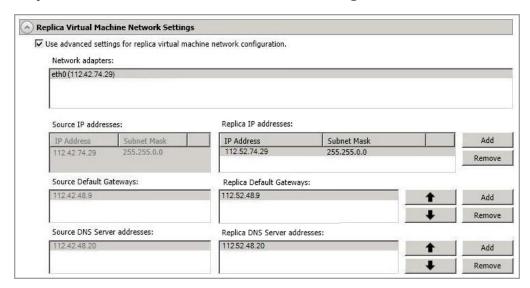
If your source has partitions, you will see them listed in the **Volume** list. Highlight a partition and set the available **Partition Properties** that are displayed to the right of the **Volume** list.



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

- **Datastore**—Specify the datastore where you want to store the .vmdk files for the partition. You can specify the location of the virtual machine configuration files on the previous **Choose Volumes to Protect** page.
- Replica disk format—Specify the format of the disk that will be created.
 - Flat—This disk format allocates the full amount of the disk space immediately, but
 does not initialize the disk space to zero until it is needed. This disk format is only
 available on ESX 5; if you select this disk type on ESX 4, a thick disk will be created.
 - **Thick**—This disk format allocates the full amount of the disk space immediately, initializing all of the allocated disk space to zero.
 - Thin—This disk format does not allocate the disk space until it is needed.
- **Disk size**—This field displays the size of the partition on the source.
- **Used space**—This field displays the amount of disk space in use on the source partition.
- **Replica volume size**—Specify the size, in MB or GB, of the replica partition on the target. The value must be at least the size of the specified **Used space** on that partition.

Replica Virtual Machine Network Settings



- Use advanced settings for replica virtual machine network configuration—Select this option to enable the replica virtual machine network setting configuration. This setting is primarily used for WAN support.
- Network adapters—Select a network adapter from the source and specify the Replica
 IP addresses, Replica Default Gateways, and Replica DNS Server addresses to be
 used after failover. If you add multiple gateways or DNS servers, you can sort them by
 using the arrow up and arrow down buttons. Repeat this step for each network adapter on
 the source.



Updates made during failover will be based on the network adapter name when protection is established. If you change that name, you will need to delete the job and re-create it so the new name will be used during failover.

If you update one of the advanced settings (IP address, gateway, or DNS server), then you must update all of them. Otherwise, the remaining items will be left blank. If you do not specify any of the advanced settings, the replica virtual machine will be assigned the same network configuration as the source.

By default, the source IP address will be included in the target IP address list as the default address. If you do not want the source IP address to be the default address on the target after failover, remove that address from the **Replica IP addresses** list.

Mirror, Verify & Orphaned Files



- Mirror all files—All protected files will be mirrored from the source to the target.
- Mirror different files—Only those protected files that are different based on date and time, size, or attributes will be mirrored from the source to the target.
 - Mirror only if the file on the source is newer than the copy on the target— Only those protected files that are newer on the source are mirrored to the target.



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

- Use block checksum for comparisons—For those files flagged as different, the
 mirroring process can perform a block checksum comparison and send only those
 blocks that are different. Make sure you always enable this option for full server to
 ESX appliance jobs.
- **Delete orphaned files**—An orphaned file is a file that exists in the replica data on the target, but does not exist in the protected data on the source. This option specifies if orphaned files should be deleted on the target during a mirror, verification, or restoration.



Orphaned file configuration is a per target configuration. All jobs to the same target will have the same orphaned file configuration.

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

Network Route



For **Send data to the target server using this route**, Double-Take will select, by default, a target route for transmissions. If desired, specify an alternate route on the target that 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. You can also select or manually enter a public IP address (which is the public IP address of the server's NAT router) if you are using a NAT environment.

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 decompresses it and then writes it to disk. You can set the level from **Minimum** to **Maximum** to suit your needs.

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 decompressed, 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.



All jobs from a single source connected to the same IP address on a target will share the same compression configuration.

Bandwidth



Bandwidth limitations are available to restrict the amount of network bandwidth used for Double-Take data transmissions. When a bandwidth limit is specified, Double-Take never exceeds that allotted amount. The bandwidth not in use by Double-Take is available for all other network traffic.



All jobs from a single source connected to the same IP address on a target will share the same bandwidth configuration.

- **Do not limit bandwidth**—Double-Take will transmit data using 100% bandwidth availability.
- Use a fixed limit—Double-Take will transmit data using a limited, fixed bandwidth. Select
 a Preset bandwidth limit rate from the common bandwidth limit values. The Bandwidth
 field will automatically update to the bytes per second value for your selected bandwidth.
 This is the maximum amount of data that will be transmitted per second. If desired, modify
 the bandwidth using a bytes per second value. The minimum limit should be 3500 bytes per
 second.
- 11. Click **Next** to continue.
- 12. Double-Take validates that your source and target are compatible. The **Summary** page displays your options and validation items.

Errors are designated by a white X inside a red circle. Warnings are designated by a black exclamation point (!) inside a yellow triangle. A successful validation is designated by a white checkmark inside a green circle. You can sort the list by the icon to see errors, warnings, or successful validations together. Click on any of the validation items to see details. You must correct any errors before you can enable protection. Depending on the error, you may be able to click **Fix** or **Fix All** and let Double-Take correct the problem for you. For those errors that Double-Take cannot correct automatically, you will need to modify the source or target to correct the error, or you can select a different target. You must revalidate the selected servers, by clicking **Recheck**, until the validation check passes without errors.

After a job is created, the results of the validation checks are logged to the job log.

13. Once your servers have passed validation and you are ready to establish protection, click **Finish**, and you will automatically be taken to the **Manage Jobs** page.

Chapter 2 Managing and controlling full server to ESX appliance jobs

Click **Manage Jobs** from the main Double-Take Console toolbar. The **Manage Jobs** page allows you to view status information about your jobs. You can also control your jobs from this page.

The jobs displayed in the right pane depend on the server group folder selected in the left pane. Every job for each server in your console session is displayed when the **Jobs on All Servers** group is selected. If you have created and populated server groups (see *Managing servers* on page 283), then only the jobs associated with the server or target servers in that server group will be displayed in the right pane.

Overview job information displayed in the top pane

The top pane displays high-level overview information about your jobs.

Column 1 (Blank)

The first blank column indicates the state of the job.

The job is in a healthy state.

1 The job is in a warning state. This icon is also displayed on any server groups that you have created that contain a job in a warning state.

The job is in an error state. This icon is also displayed on any server groups that you have created that contain a job in an error state.

The job is in an unknown state.

Job

The name of the job

Source Server

The name of the source. This could be the name or IP address of your source.

Target Server

The name of the target. This could be the name or IP address of your target.

Job Type

Each job type has a unique job type name. This job is a Full Server to ESX Appliance job. For a complete list of all job type names, press F1 to view the Double-Take Console online help.

Activity

There are many different **Activity** messages that keep you informed of the job activity. Most of the activity messages are informational and do not require any administrator interaction. If you see error messages, check the job details. Keep in mind that **Idle** indicates console to server activity is idle, not that your servers are idle.

Mirror Status

- Calculating—The amount of data to be mirrored is being calculated.
- In Progress—Data is currently being mirrored.
- Waiting—Mirroring is complete, but data is still being written to the target.
- Idle—Data is not being mirrored.
- Paused—Mirroring has been paused.
- Stopped—Mirroring has been stopped.
- Removing Orphans—Orphan files on the target are being removed or deleted depending on the configuration.
- Verifying—Data is being verified between the source and target.
- Unknown—The console cannot determine the status.

Replication Status

- Replicating—Data is being replicated to the target.
- Ready—There is no data to replicate.
- Pending—Replication is pending.
- Stopped—Replication has been stopped.
- Out of Memory—Replication memory has been exhausted.
- Failed—The Double-Take service is not receiving replication operations from the Double-Take driver. Check the Event Viewer for driver related issues.
- Unknown—The console cannot determine the status.

Transmit Mode

- Active—Data is being transmitted to the target.
- Paused—Data transmission has been paused.
- Scheduled—Data transmission is waiting on schedule criteria.
- Stopped—Data is not being transmitted to the target.
- Error—There is a transmission error.
- Unknown—The console cannot determine the status.

Detailed job information displayed in the bottom pane

The details displayed in the bottom pane of the **Manage Jobs** page provide additional information for the job highlighted in the top pane. If you select multiple jobs, the details for the first selected job will be displayed.

Name

The name of the job

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.
- **Busy**—The source is low on memory causing a delay in getting the state of the data on the target.
- **Not Loaded**—Double-Take target functionality is not loaded on the target server. This may be caused by an activation code error.
- Not Ready—The Linux drivers have not yet completed loading on the target.
- Unknown—The console cannot determine the status.

Mirror remaining

The total number of mirror bytes that are remaining to be sent from the source to the target

Mirror skipped

The total number of bytes that have been skipped when performing a difference or checksum mirror. These bytes are skipped because the data is not different on the source and target.

Replication queue

The total number of replication bytes in the source queue

Disk queue

The amount of disk space being used to gueue data on the source

Bytes sent

The total number of mirror and replication bytes that have been transmitted to the target

Bytes sent (compressed)

The total number of compressed mirror and replication bytes that have been transmitted to the target. If compression is disabled, this statistic will be the same as **Bytes sent**.

Connected since

The date and time indicating when the current job was made. This field is blank, indicating that a TCP/IP socket is not present, when the job 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.

Recent activity

Displays the most recent activity for the selected job, along with an icon indicating the success or failure of the last initiated activity. Click the link to see a list of recent activities for the selected job. You can highlight an activity in the list to display additional details about the activity.

Additional information

Depending on the current state of your job, you may see additional information displayed to keep you informed about the progress and status of your job. If there is no additional information, you will see (None) displayed.

Job controls

You can control your job through the toolbar buttons available on the **Manage jobs** page. If you select multiple jobs, some of the controls will apply only to the first selected job, while others will apply to all of the selected jobs. For example, View Job Details will only show details for the first selected job, while **Stop** will stop protection for all of the selected jobs.

If you want to control just one job, you can also right click on that job and access the controls from the pop-up menu.

Create a New Joh



This button leaves the **Manage Jobs** page and opens the **Get Started** page.

View Job Details



This button leaves the **Manage Jobs** page and opens the **View Job Details** page.

Delete III



Stops (if running) and deletes the selected jobs.

If you no longer want to protect the source and no longer need the replica of the source on the target, select to delete the associated replica virtual machine. Selecting this option will remove the job and completely delete the replica virtual machine on the target.

If you no longer want to mirror and replicate data from the source to the target but still want to keep the replica of the source on the target, select to keep the associated replica virtual machine.

Provide Credentials



Changes the login credentials that the job (which is on the target machine) uses to authenticate to the servers in the job. This button opens the Provide Credentials dialog box where you can specify the new account information and which servers you want to update. See Providing server credentials on page 287. You will remain on the Manage **Jobs** page after updating the server credentials. If your servers use the same credentials, make sure you also update the credentials on the Manage Servers page so that the Double-Take Console can authenticate to the servers in the console session. See *Managing servers* on page 283.

View Recent Activity



Displays the recent activity list for the selected job. Highlight an activity in the list to display additional details about the activity.



Starts or resumes the selected jobs.

If you have previously stopped protection, the job will restart mirroring and replication.

If you have previously paused protection, the job will continue mirroring and replication from where it left off, as long as the Double-Take queue was not exhausted during the time the job was paused. If the Double-Take queue was exhausted during the time the job was paused, the job will restart mirroring and replication.

Also if you have previously paused protection, all jobs from the same source to the same IP address on the target will be resumed.



Pauses the selected jobs. Data will be queued on the source while the job is paused.

All jobs from the same source to the same IP address on the target will be paused.



Stops the selected jobs. The jobs remain available in the console, but there will be no mirroring or replication data transmitted from the source to the target. Mirroring and replication data will not be gueued on the source while the job is stopped, requiring a remirror when the job is restarted. The type of remirror will depend on your job settings.

Take Snapshot



Snapshots are not applicable to full server to ESX appliance jobs.

Manage Snapshots



Snapshots are not applicable to full server to ESX appliance jobs.

Failover or Cutover



Starts the failover process. See Failing over full server to ESX appliance jobs on page 274 for the process and details of failing over a full server to ESX appliance job.

Failback



Starts the failback process. Failback does not apply to full server to ESX appliance jobs.

Restore 🚨



Starts the restoration process. Restore does not apply to full server to ESX appliance jobs.



Reverses protection. Reverse protection does not apply to full server to ESX appliance jobs.

Undo Failover



Cancels a test failover by undoing it. Undo failover does not apply to full server to ESX appliance jobs.

View Job Log



Opens the job log. On the right-click menu, this option is called **View Logs**, and you have the option of opening the job log, source server log, or target server log. See Viewing the log files through the Double-Take Console on page 279 for details on all three of these logs.

Other Job Actions



Opens a small menu of other job actions. These job actions are not available for full server to ESX appliance jobs.

Filter

Select a filter option from the drop-down list to only display certain jobs. You can display Healthy jobs, Jobs with warnings, or Jobs with errors. To clear the filter, select All jobs. If you have created and populated server groups, then the filter will only apply to the jobs associated with the server or target servers in that server group. See Managing servers on page 283.

Type a server name

Displays only jobs that contain the text you entered. If you have created and populated server groups, then only jobs that contain the text you entered associated with the server or target servers in that server group will be displayed. See Managing servers on page 283.

Overflow Chevron



Displays any toolbar buttons that are hidden from view when the window size is reduced.

Viewing full server to ESX appliance job details

From the **Manage Jobs** page, highlight the job and click **View Job Details** in the toolbar.

Review the following table to understand the detailed information about your job displayed on the **View Job Details** page.

Job name

The name of the job

Job type

Each job type has a unique job type name. This job is a Full Server to ESX Appliance job. For a complete list of all job type names, press F1 to view the Double-Take Console online help.

Health

- The job is in a healthy state.
- 1 The job is in a warning state.
- The job is in an error state.
- The job is in an unknown state.

Activity

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

Connection ID

The incremental counter used to number connections. The number is incremented when a connection is created. It is also incremented by internal actions, such as an auto-disconnect and auto-reconnect. The lowest available number (as connections are created, stopped, deleted, and so on) will always be used. The counter is reset to one each time the Double-Take service is restarted.

Transmit mode

- Active—Data is being transmitted to the target.
- Paused—Data transmission has been paused.
- Scheduled—Data transmission is waiting on schedule criteria.
- Stopped—Data is not being transmitted to the target.
- **Error**—There is a transmission error.
- Unknown—The console cannot determine the status.

Target data state

- OK—The data on the target is in a good state.
- **Mirroring**—The target is in the middle of a mirror process. The data will not be in a good state until the mirror is complete.
- Mirror Required—The data on the target is not in a good state because a remirror is required. This may be caused by an incomplete or stopped mirror or an operation may have been dropped on the target.
- Busy—The source is low on memory causing a delay in getting the state of the data on the target.
- Not Loaded—Double-Take target functionality is not loaded on the target server.
 This may be caused by an activation code error.
- Not Ready—The Linux drivers have not yet completed loading on the target.
- Unknown—The console cannot determine the status.

Target route

The IP address on the target used for Double-Take transmissions.

Compression

- On / Level—Data is compressed at the level specified.
- Off—Data is not compressed.

Encryption

- On—Data is being encrypted before it is sent from the source to the target.
- Off—Data is not being encrypted before it is sent from the source to the target.

Bandwidth limit

If bandwidth limiting has been set, this statistic identifies the limit. The keyword **Unlimited** means there is no bandwidth limit set for the job.

Connected since

The date and time indicating when the current job was made. This field is blank, indicating that a TCP/IP socket is not present, when the job 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.

Additional information

Depending on the current state of your job, you may see additional information displayed to keep you informed about the progress and status of your job. If there is no additional information, you will see (None) displayed.

Mirror status

- Calculating—The amount of data to be mirrored is being calculated.
- In Progress—Data is currently being mirrored.
- Waiting—Mirroring is complete, but data is still being written to the target.

- Idle—Data is not being mirrored.
- Paused—Mirroring has been paused.
- Stopped—Mirroring has been stopped.
- Removing Orphans—Orphan files on the target are being removed or deleted depending on the configuration.
- Verifying—Data is being verified between the source and target.
- Unknown—The console cannot determine the status.

Mirror percent complete

The percentage of the mirror that has been completed

Mirror remaining

The total number of mirror bytes that are remaining to be sent from the source to the target

Mirror skipped

The total number of bytes that have been skipped when performing a difference or checksum mirror. These bytes are skipped because the data is not different on the source and target.

Replication status

- Replicating—Data is being replicated to the target.
- Ready—There is no data to replicate.
- Pending—Replication is pending.
- Stopped—Replication has been stopped.
- Out of Memory—Replication memory has been exhausted.
- Failed—The Double-Take service is not receiving replication operations from the Double-Take driver. Check the Event Viewer for driver related issues.
- Unknown—The console cannot determine the status.

Replication queue

The total number of replication bytes in the source queue

Disk queue

The amount of disk space being used to queue data on the source

Bytes sent

The total number of mirror and replication bytes that have been transmitted to the target

Bytes sent compressed

The total number of compressed mirror and replication bytes that have been transmitted to the target. If compression is disabled, this statistic will be the same as **Bytes sent**.

Validating a full server to ESX appliance job

Over time, you may want to confirm that any changes in your network or environment have not impacted your Double-Take job. Use these instructions to validate an existing job.

- 1. From the Manage Jobs page, highlight the job and click View Job Details in the toolbar.
- 2. In the **Tasks** area on the right on the **View Job Details** page, click **Validate job properties**.
- Double-Take validates that your source and target are compatible. The Summary page displays your options and validation items.
 - Errors are designated by a white X inside a red circle. Warnings are designated by a black exclamation point (!) inside a yellow triangle. A successful validation is designated by a white checkmark inside a green circle. You can sort the list by the icon to see errors, warnings, or successful validations together. Click on any of the validation items to see details. You must correct any errors before you can enable protection. Depending on the error, you may be able to click **Fix** or **Fix All** and let Double-Take correct the problem for you. For those errors that Double-Take cannot correct automatically, you will need to modify the source or target to correct the error, or you can select a different target. You must revalidate the selected servers, by clicking **Recheck**, until the validation check passes without errors.
- 4. Once your servers have passed validation, click Close.

Editing a full server to ESX appliance job

Use these instructions to edit a full server to ESX appliance job.

- 1. From the **Manage Jobs** page, highlight the job and click **View Job Details** in the toolbar.
- 2. In the **Tasks** area on the right on the **View Job Details** page, click **Edit job properties**.
- 3. You will see the same options for your full server to ESX appliance job as when you created the job, but you will not be able to edit all of them. If desired, edit those options that are configurable for an existing job. See *Creating a full server to ESX appliance job* on page 236 for details on each job option.



Changing some options may require Double-Take to automatically disconnect, reconnect, and remirror the job.

4. If you want to modify the workload items or replication rules for the job, click **Edit workload or replication rules**. Modify the **Workload item** you are protecting, if desired. Additionally, you can modify the specific **Replication Rules** for your job.

Volumes and folders with a green highlight are included completely. Volumes and folders highlighted in light yellow are included partially, with individual files or folders included. If there is no highlight, no part of the volume or folder is included. To modify the items selected, highlight a volume, folder, or file and click **Add Rule**. Specify if you want to **Include** or **Exclude** the item. Also, specify if you want the rule to be recursive, which indicates the rule should automatically be applied to the subdirectories of the specified path. If you do not select **Recursive**, the rule will not be applied to subdirectories.

You can also enter wildcard rules, however you should do so carefully. Rules are applied to files that are closest in the directory tree to them. If you have rules that include multiple folders, an exclusion rule with a wild card will need to be added for each folder that it needs applied to. For example, if you want to exclude all .log files from D:\ and your rules include D:\, D:\Dir1, and D:\Dir2, you would need to add the exclusion rule for the root and each subfolder rule. So you will need to add exclude rules for D:*.log, D:\Dir1*.log, and D:\Dir2*.log.

If you need to remove a rule, highlight it in the list at the bottom and click **Remove Rule**. Be careful when removing rules. Double-Take may create multiple rules when you are adding directories. For example, if you add E:\Data to be included in protection, then E:\ will be excluded. If you remove the E:\ exclusion rule, then the E:\Data rule will be removed also.

Click **OK** to return to the **Edit Job Properties** page.

- 5. Click **Next** to continue.
- 6. Double-Take validates that your source and target are compatible. The **Summary** page displays your options and validation items.

Errors are designated by a white X inside a red circle. Warnings are designated by a black exclamation point (!) inside a yellow triangle. A successful validation is designated by a white checkmark inside a green circle. You can sort the list by the icon to see errors, warnings, or successful validations together. Click on any of the validation items to see details. You must correct any errors before you can enable protection. Depending on the error, you may be able to

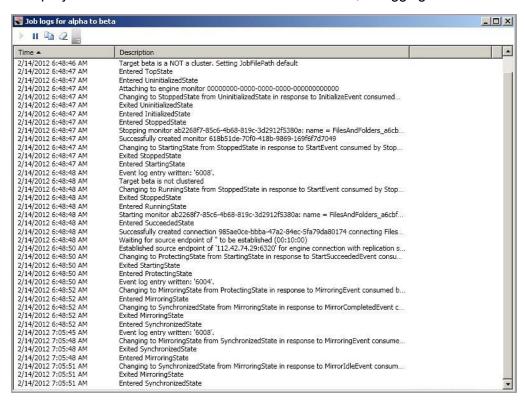
click **Fix** or **Fix All** and let Double-Take correct the problem for you. For those errors that Double-Take cannot correct automatically, you will need to modify the source or target to correct the error, or you can select a different target. You must revalidate the selected servers, by clicking **Recheck**, until the validation check passes without errors.

After a job is created, the results of the validation checks are logged to the job log.

7. Once your servers have passed validation and you are ready to update your job, click **Finish**.

Viewing a full server to ESX appliance job log

You can view a job log file through the Double-Take Console by selecting **View Job Log** from the toolbar on the **Manage Jobs** page. Separate logging windows allow you to continue working in the Double-Take Console while monitoring log messages. You can open multiple logging windows for multiple jobs. When the Double-Take Console is closed, all logging windows will automatically close.



The following table identifies the controls and the table columns in the **Job logs** window.



This button starts the addition and scrolling of new messages in the window.



This button pauses the addition and scrolling of new messages in the window. This is only for the **Job logs** window. The messages are still logged to their respective files on the server.

Copy 🍱

This button copies the messages selected in the **Job logs** window to the Windows clipboard.

Clear 2

This button clears the **Job logs** window. The messages are not cleared from the respective files on the server. If you want to view all of the messages again, close and reopen the **Job logs** window.

Time

This column in the table indicates the date and time when the message was logged.

Description

This column in the table displays the actual message that was logged.

Statistics

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 statistics file configuration are detected and applied immediately without restarting the Double-Take service.

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

Sample DTStat output

```
0/11/10 12:48:05:2040
SYSTEMALLOCATOR::Total Bytes: 0
IQALLOCATOR::Total Bytes: 0
SECURITY::Logins : 1 FailedLogins : 0
KERNEL::SourceState: 2 TargetState: 1 Start Time: Tue Sep 11 12:45:26 2007
RepOpsGenerated: 436845 RepBytesGenerated: 0
MirOpsGenerated: 3316423 MirBytesGenerated: 108352749214952
   FailedMirrorCount: 0 FailedRepCount: 0
   ActFailCount: 0 TargetOpenHandles: 0 DriverQueuePercent: 0
TARGET:: PeerAddress: 10.10.1.104 LocalAddress: 10.10.1.104
   Ops Received: 25 Mirror Ops Received: 23
   Retries: 0 OpsDropped: 0 Ops Remaining: 0
Orphan Files Removed: 0 Orphan Directories Removed: 0 Orphan Bytes Removed: 0
   Bytes In Target Queue: 0 Bytes In Target Disk Queue: 0 TasksSucceeded: 0 TasksFailed: 0 TasksIgnored: 0
SOURCE::autoDisConnects: 0 autoReConnects: 1
   lastFileTouched : /log/data file
CONNECTION:: conPeerAddress: 10.10.1.104
   connectTime: Tue Sep 11 12:45:34 2007
   conState: 1 conOpsInCmdQueue: 0 conOpsInAckQueue: 0
   conOpsInRepQueue: 0 conOpsInMirQueue: 0 conBytesInRepQueue: 0
   conOpsTx: 27 conBytesInMirQueue: 0 conBytesTx: 14952687269
   conBytesCompressedTx: 14952
   conOpsRx: 201127 conBytesRx: 647062280 conResentOpCount: 0 conBytesInDiskQueue: 0
   conBandwidthLimit: 429496295 conBytesSkipped: 22867624 conMirrorBytesRemain: 0
   conMirrorPercent: 100.0%
   conTaskCmdsSubmitted: 0 conTaskCmdsQueued: 0
   conTasksSucceeded: 0 conTasksFailed: 0 conTasksIgnored: 0
```

Viewing the statistics file

The statistics log file created is a binary file. To view the log file, you must run the DTStat utility from a command prompt. From the directory where Double-Take is installed, run the DTStat command.

Command

DTSTAT

Description

Starts the DTStats statistics logging utility from a command prompt

Syntax

DTSTAT [-p][-i < interval>][-t < filename>] [-f < filename>] [-s < filename>] [-st < filename>][-IP < address>] [-START < mm/dd/yyyyhh:mm>][-STOP < mm/dd/yyyyhh:mm>] [-SERVER < ip_address> < port_number>]

Options

- -p—Do not print the output to the screen. This option will increase the speed of the output to files.
- -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/yyyyhh:mm—Filters out any data prior to the specified date and time
- -STOP mm/dd/yyyyhh: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 -p -f statistic.sts -s statistic.csv
- DTStat -p -f statistic.sts -st statistic.csv

Notes

- This command is not case-sensitive.
- If no options are specified, DTStat will print the output to the screen at an interval of every one second.
- If the statistics are not changing, DTStat will discontinue writing until statistics begin updating again.

Statistics

The following table identifies the Double-Take statistics.



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

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

Statistic values are cumulative. For example if Kernel, RepBytesGenerated is 10000 at 1:00pm and 25000 at 2:00pm, the difference is 15000 and that is the amount of change that occurred within that one hour.

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

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

Date/Time Stamp

The date and time that the snapshot was taken. This is the date and time that each statistic was logged. By default, these are generated once a second, as long as there are statistics being generated. If mirroring/replication is idle, then DTStat will be idle as well.

System Allocator, Total Bytes

The number of bytes currently allocated to the system pagefile

IQAllocator, Total Bytes

The number of bytes currently allocated to the intermediate queue

Security, Logins

The number of successful login attempts

Security, Failed Logins

The number of failed login attempts

Kernel, SourceState

- 0—Source is not running
- 1—Source is running without the replication driver

2—Source is running with the replication driver

Kernel, TargetState

- 0—Target is not running
- 1—Target is running

Kernel, Start Time

Date and time stamp indicating when the Double-Take service was loaded

Kernel, RepOpsGenerated

The number of replication operations generated by the file system driver. An op is a file system operation. Double-Take replicates data by sending the file system operations across the network to the target. RepOpsGenerated indicates the number of file system operations that have been generated by replication.

Kernel, RepBytesGenerated

The number of replication bytes generated by the file system driver. This is the number of bytes generated during replication. In other words, this is roughly the amount of traffic being sent across the network that is generated by replication. It does not take into account TCP/IP overhead (headers and such).

Kernel, MirOpsGenerated

The number of mirror operations transmitted to the target. Mirroring is completed by transmitting the file system operations necessary to generate the files on the target. This statistic indicates the number of file system operations that were transmitted during the initial mirror. It will continue to increase until the mirror is complete. Any subsequent remirrors will reset this field to zero and increment from there.

Kernel, MirBytesGenerated

The number of mirror bytes transmitted to the target. This is the number of bytes generated during mirroring. In other words, this is roughly the amount of traffic being sent across the network that is generated by the mirror. It does not take into account TCP/IP overhead (headers and such), however it does account for attributes and other overhead associated with creating a file. With many small files in a directory, you will see larger statistics than expected because of the file creation overhead. Any subsequent remirror will reset this field to zero and increment from there.

Kernel, FailedMirrorCount

The number of mirror operations that failed due to an error reading the file from the disk

Kernel, FailedRepCount

The number of replication operations that failed due to an error reading the file from the disk

Kernel, ActFailCount

The number of activation code failures when loading the source or target. Activation codes can be bad for reasons such as: expiration of evaluation codes, duplicate codes,

incorrect codes, etc.

Kernel, TargetOpenHandles

The number of handles currently open on the target

Kernel, DriverQueuePercent

The amount of throttling calculated as a percentage of the stop replicating limit

Target, PeerAddress

The IP address of the source machine

Target, LocalAddress

The IP address of the target machine.

Target, Ops Received

The total number of operations received by this machine as a target since the Double-Take service was loaded

Target, Mirror Ops Received

The total number of mirror operations received by this machine as a target since the Double-Take service was loaded. This number does not reset to zero for remirrors.

Target, Retries

The number of retries performed before all operations were completed

Target, OpsDropped

The number of operations skipped during a difference mirror. During a difference mirror, if Double-Take detects that there have been no changes to a file, then it will indicate the number of operations it did not send for this file in this field.

Target, Ops Remaining

The total number of operations that are left in the target queue

Target, Orphan Files Removed

The number of orphan files removed from the target machine

Target, Orphan Directories Removed

The number of orphan directories removed from the target machine

Target, Orphan Bytes Removed

The number of orphan bytes removed from the target machine

Target, Bytes In Target Queue

The number of bytes currently in the system memory queue on the target

Target. Bytes In Target Disk Queue

The number of bytes currently in the disk queue on the target

Target, TasksSucceeded

The number of task commands that have succeeded on the target

Target, TasksFailed

The number of task commands that have failed on the target

Target, TasksIgnored

The number of task commands that have been ignored on the target

Source, autoDisConnects

The number of automatic disconnects since starting Double-Take. Auto-disconnects occur because the source no longer sees the target This could be because the connection between the two has failed at some point or because the target machine data is changing on the source faster than the source can get the data to the target. This field tracks the number of times an auto-disconnect has occurred since the Double-Take service was started.

Source, autoReConnects

The number of automatic reconnects since starting Double-Take. Auto-reconnect occurs after a target machine is back online. This field tracks the number of times an auto-reconnect has happened since the Double-Take service was started.

Source, lastFileTouched

The last filename that had a replication operation executed

Connection, conPeerAddress

The IP address of the target machine

Connection, connectTime

The time that this connection was established

Connection, conState

The state of the active connection

- 0—None. This indicates there is no active connection. This may be because the
 connection has not been established or the underlying connection is unavailable.
 Statistics are still available for the source and target machines.
- 1—Active. This indicates that the connection is functioning normally and has no scheduling restrictions imposed on it at this time. (There may be restrictions, but it is currently in a state that allows it to transmit.)
- 2—Paused. This indicates a connection that has been paused.
- 4—Scheduled. This indicates a connection that is not currently transmitting due to scheduling restrictions (bandwidth limitations, time frame limitations, and so on).

• 8—Error. This indicates a connection that is not transmitting because something has gone wrong (for example, lost connection).

Only the Scheduled and Error states can coexist. All other states are mutually exclusive. Statistics will display a conState of 12 when the connection is in both a scheduled and an error state because this is the sum of the two values (4 + 8).

Connection, conOpsInCmdQueue

The number of operations waiting to be executed on the target

Connection, conOpsInAckQueue

The number of operations waiting in the acknowledgement queue. Each operation that is generated receives an acknowledgement from the target after that operation has been received by the target. This statistic indicates the number of operations that have yet to receive acknowledgement of receipt.

Connection, conOpsInRepQueue

The number of replication operations currently waiting to be executed on the target

Connection, conOpsInMirQueue

The number of mirror operations currently waiting to be executed on the target

Connection, conBytesInRepQueue

The number of replication bytes remaining to be transmitted to the target

Connection, conOpsTx

The number of operations transmitted to the target. This is the total number of operations that Double-Take has transmitted as a source. In other words, the cumulative number of operations transmitted by this source to all connected targets.

Connection, conBytesInMirQueue

The number of mirror bytes remaining to be transmitted to the target

Connection, conBytesTx

The number of bytes transmitted to the target. This is the total number of bytes that Double-Take has transmitted as a source. In other words, the cumulative number of bytes transmitted by this source to all connected targets.

Connection, conBytesCompressedTx

The number of compressed bytes transmitted to the target.

Connection, conOpsRx

The number of operations received by the target. The number of operations that the target for this connection (as indicated by the IP address field) has received from this source.

Connection, conBytesRx

The number of bytes received by the target. The number of bytes that the target for this connection (as indicated by the IP address field) has received from this source.

Connection, conResentOpCount

The number of operations resent because they were not acknowledged

Connection, conBytesInDiskQueue

The number of bytes in the source disk queue

Connection, conBandwidthLimit

The amount of bandwidth that may be used to transfer data

Connection, conBytesSkipped

The number of bytes skipped during a difference mirror. During a difference mirror, if Double-Take detects that there have been no changes to a file, then it will indicate the number of bytes it did not send for this file in this field.

Connection, conMirrorBytesRemaining

The number of mirror bytes remaining to be transmitted

Connection, conMirrorPercent

The percentage of the mirror that has been completed. This field is determined if the Job size was calculated.

Connection, conTaskCmdsSubmitted

The number of task commands that have been submitted on the source

Connection, conTaskCmdsQueued

The number of task commands that have been queued on the source

Connection, conTasksSucceeded

The number of task commands that have succeeded on the source

Connection, conTasksFailed

The number of task commands that have failed on the source

Connection, conTasksIgnored

The number of task commands that have been ignored on the source

Chapter 3 Failing over full server to ESX appliance jobs

You will be notified in the console when a failover condition has been met. At this time, you should trigger failover. You can also trigger failover at any other time you desire, thus allowing you to better control the failover process.

- 1. On the **Manage Jobs** page, highlight the job that you want to failover and click **Failover or Cutover** in the toolbar.
- 2. Select the type of failover to perform.
 - Failover to live data—Select this option to initiate a full, live failover using the current data
 on the target. This option will shutdown the source machine (if it is online), stop the
 protection job, and start the replica virtual machine on the target with full network
 connectivity.
 - Failover to a snapshot—This option is not available for full server to ESX appliance jobs.
- 3. Select how you want to handle the data in the target queue. You may want to check the amount of data in queue on the target by reviewing the *Statistics* on page 265.
 - Discard data in the target queues and failover or cutover immediately—All of the data in the target queue will be discarded and failover will begin immediately. The advantage to this option is that failover will occur immediately. The disadvantage is that any data in the target queue will be lost.
- 4. When you are ready to begin failover, click Failover.



If you need to update DNS after failover, there is a sample DNS update script located in /etc/DT/sysprep.d. You may need to modify the script for your environment. If you need basic assistance with script modifications, contact Technical Support. Assistance with advanced scripting will be referred to Professional Services.

There is no reverse or failback once you have failed over.

Chapter 4 Double-Take Console

Each time you open the Double-Take Console, you start at the **Home** page. This page provides a high-level overview of the status of your jobs.

The appearance of the **Home** page is the same for all users. However, other console pages may have variances in the appearance depending on the Double-Take products that you have installed, the Double-Take activation codes on your servers, and the type of job you are working with.

- **Headlines**—The top section gives a quick overview of any jobs that require attention as well as providing quick access buttons.
 - These jobs require attention—Any jobs that require attention (those in an error state) are listed. You will see the source and target server names listed, as well as a short description of the issue that requires your attention. If the list is blank, there are no jobs that require immediate attention.
 - Tools—Select this drop-down list to launch other Vision Solutions consoles.
- Servers Summary—The middle section summarizes the servers in your console.
 - Total number of servers—This field displays the number of servers that you have been added to the console.
- Jobs Summary—The bottom section summarizes the jobs in your console.
 - **Total number of jobs**—This field displays the number of jobs running on the servers in your console.

Double-Take Console requirements

You must meet the following requirements for the Double-Take Console.

- Operating system—The Double-Take Console can be run from a Windows source or target. It can also be run from a 32-bit or 64-bit physical or virtual machine running Windows 8, Windows 7, Windows Vista, or Windows XP Service Pack 2 or later.
- Microsoft .NET Framework
 — Microsoft .NET Framework version 3.5 Service Pack 1 is required. This version is not included in the .NET version 4.0 release. Therefore, even if you have .NET version 4.0 installed, you will also need version 3.5.1. For Windows 2008 and earlier, you can install this version from the Double-Take DVD, via a web connection during the Double-Take installation, or from a copy you have obtained manually from the Microsoft web site. For Windows 2008 R2 and later, you need to enable it through Windows features.
- Screen resolution—For best results, use a 1024x768 or higher screen resolution.



The Double-Take installation prohibits the console from being installed on Server Core. Because Windows 2012 allows you to switch back and forth between Server Core and a full installation, you may have the console files available on Server Core, if you installed Double-Take while running in full operating system mode. In any case, you cannot run the Double-Take Console on Server Core.

Console options

There are several options that you can set that are specific to the Double-Take Console.

- Monitoring interval—Specifies how often, in seconds, the console refreshes the monitoring data. The servers will be polled at the specified interval for information to refresh the console.
- **Automatic retry**—This option will have the console automatically retry server login credentials, after the specified retry interval, if the server login credentials are not accepted. Keep in mind the following caveats when using this option.
 - This is only for server credentials, not job credentials.
 - A set of credentials provided for or used by multiple servers will not be retried for the specified retry interval on any server if it fails on any of the servers using it.
 - Verify your environment's security policy when using this option. Check your policies for failed login lock outs and resets. For example, if your policy is to reset the failed login attempt count after 30 minutes, set this auto-retry option to the same or a slightly larger value as the 30 minute security policy to decrease the chance of a lockout.
 - Restarting the Double-Take Console will automatically initiate an immediate login.
 - Entering new credentials will initiate an immediate login using the new credentials.
- **Retry on this interval**—If you have enabled the automatic retry, specify the length of time, in minutes, to retry the login.
- **Default port for XML web services protocol**—Specifies the port that the console will use when sending and receiving data to Double-Take servers. By default, the port is 6325. Changes to the console port will not take effect until the console is restarted.
- **Default port for legacy protocol**—If you are using an older Double-Take version, you will need to use the legacy protocol port. This applies to Double-Take versions 5.1 or earlier.
- Export Diagnostic Data—This button creates a raw data file that can be used for debugging errors in the Double-Take Console. Use this button as directed by technical support.
- View Log File—This button opens the Double-Take Console log file. Use this button as directed by technical support.
- View Data File—This button opens the Double-Take Console data file. Use this button as
 directed by technical support.
- Automatically check for updates—By default, each time the console is started, it will
 automatically check the Vision Solutions web site to see if there is updated console software
 available. If there is updated console software available, an Automatic Updates section will
 appear on the Home page. Click Get the latest update to download and install the updated
 console software.

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

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

- Update available—If there is an update available, click Get Update. The dialog box will
 close and your web browser will open to the Vision Solutions web site where you can
 download and install the update.
- No update available—If you are using the most recent console software, that will be indicated. Click Close.

- No connection available—If the console cannot contact the update server of if there is an
 error, the console will report that information. The console log contains a more detailed
 explanation of the error. Click Check using Browser if you want to open your browser to
 check for console software updates. You will need to use your browser if your Internet
 access is through a proxy server.
- **Enable license inventory**—This option allows you to use this console to manage the Double-Take licenses assigned to your organization. When this option is enabled, the **Manage License Inventory** page is also enabled.



The license inventory feature may not appear in your console if your service provider has restricted access to it.

Viewing the log files through the Double-Take Console

You can view the Double-Take, Double-Take Management Service, job, and Double-Take Console log files through the Double-Take Console.

- Viewing the Double-Take and Double-Take Management Service logs on page 279
- Viewing the job log file on page 281
- Viewing the Double-Take Console log file on page 282

Viewing the Double-Take and Double-Take Management Service logs

You can view the Double-Take and Double-Take Management Service log files through the Double-Take Console using either of these two methods.

- On the Manage Servers page, highlight a server in the list and click View Server Logs from the toolbar.
- On the Manage Jobs page, right-click a job and select View Logs. Select either the source server log or the target server log.

Separate logging windows allow you to continue working in the Double-Take Console while monitoring log messages. You can open multiple logging windows for multiple servers. When the Double-Take Console is closed, all logging windows will automatically close.

The following table identifies the controls and the table columns in the **Server logs** window.



This button starts the addition and scrolling of new messages in the window.



This button pauses the addition and scrolling of new messages in the window. This is only for the **Server logs** window. The messages are still logged to their respective files on the server.



This button copies the messages selected in the **Server logs** window to the Windows clipboard.

Clear 2

This button clears the **Server logs** window. The messages are not cleared from the respective files on the server. If you want to view all of the messages again, close and reopen the **Server logs** window.

Filter

From the drop-down list, you can select to view all log messages or only those messages from the Double-Take engine or the Double-Take Management Service.

Time

This column in the table indicates the date and time when the message was logged.

Description

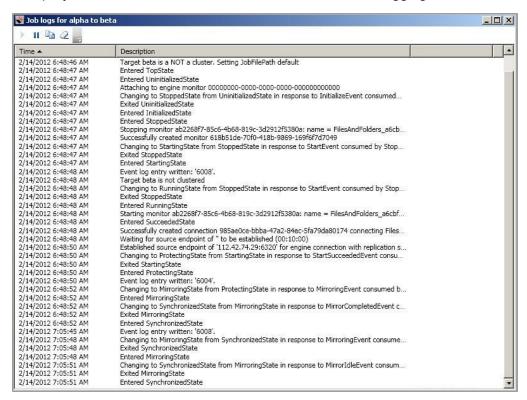
This column in the table displays the actual message that was logged.

Service

This column in the table indicates if the message is from the Double-Take engine or the Double-Take Management Service.

Viewing the job log file

You can view a job log file through the Double-Take Console by selecting **View Job Log** from the toolbar on the **Manage Jobs** page. Separate logging windows allow you to continue working in the Double-Take Console while monitoring log messages. You can open multiple logging windows for multiple jobs. When the Double-Take Console is closed, all logging windows will automatically close.



The following table identifies the controls and the table columns in the **Job logs** window.



This button starts the addition and scrolling of new messages in the window.



This button pauses the addition and scrolling of new messages in the window. This is only for the **Job logs** window. The messages are still logged to their respective files on the server.

Сору

This button copies the messages selected in the **Job logs** window to the Windows clipboard.

Clear 4

This button clears the **Job logs** window. The messages are not cleared from the respective files on the server. If you want to view all of the messages again, close and reopen the **Job logs** window.

Time

This column in the table indicates the date and time when the message was logged.

Description

This column in the table displays the actual message that was logged.

Viewing the Double-Take Console log file

You can view the Double-Take Console log file by using the following instructions.

- 1. Select **Tools**, **Options**.
- 2. Expand the **Diagnostics** section, if necessary.
- 3. Click View Log File.
- 4. Click **Cancel** to return back to the console page you were on previously. (Unless you were in the middle of a job creation workflow, in which case you will be returned to the beginning of the workflow.)

The log file is opened in the default text editor. The file will remain open until you close.

Chapter 5 Managing servers

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

You can also organize the servers that are in your console into groups, allowing you to filter the servers you are viewing based on your organization. The servers displayed in the right pane depend on the server group folder selected in the left pane. Every server in your console session is displayed when the **All Servers** group is selected. If you have created and populated server groups under **My Servers**, then only the servers in the selected group will displayed in the right pane.



If you have uninstalled and reinstalled Double-Take on a server, you may see the server twice on the **Manage Servers** page because the reinstall assigns a new unique identifier to the server. One of the servers (the original version) will show with the red X icon. You can safely remove that server from the console.

Right pane display

The following table identifies the columns displayed in the right pane of the **Manage Servers** page.

Column 1 (Blank)

The first blank column indicates the machine type.

- Double-Take source or target server which could be a physical server, virtual machine, or a cluster node
- Double-Take source or target server which is a Windows cluster
- SVMware server which could be a vCenter server or an ESX or ESXi host.
- Double-Take controller appliance
- 🗟 Double-Take replication appliance
- Double-Take Reporting Service server
- Offline server which means the console cannot communicate with this machine.
- Server error which means the console can communicate with the machine, but it cannot communicate with Double-Take on it.

Column 2 (Blank)

The second blank column indicates the security level

- Processing—The console is attempting to communicate with machine.
- Administrator access—This level grants full control.
- Monitor only access—This level grants monitoring privileges only.
- No security access—This level does not allow monitoring or control.

Server

The name or IP address of the server. If you have specified a reserved IP address, it will be displayed in parenthesis.

Activity

There are many different **Activity** messages that keep you informed of the server activity. Most of the activity messages are informational and do not require any administrator interaction. If you see error messages, check the server details. See *Viewing server details* on page 288.

Version

The product version information

Licensing Status

The status of the license on the server. If your license is expired, any jobs using that server will be in an error state.

Product

The Double-Take products licensed for the server or the Double-Take role for the server.

Activation Code

The activation codes associated with the products licensed for the server. If your license is not valid for the operating system on your server, the activation code will be identified as Invalid Activation Code. There will be no activation code listed for those servers that are not licensed, like a VMware server.

Serial Number

The serial number associated with the activation code

Main toolbar and right-click menu

The following options are available on the main toolbar of the **Manage Servers** page and the right-click menu. Some of the options are only available in the right-click menu. Each of the options control the server that is selected in the right pane.

Add Servers



Adds a new server. This button leaves the **Manage Servers** page and opens the **Add Servers** page. See *Adding a virtual recovery appliance to the console* on page 225 and *Adding a Linux source to the console* on page 233.

Add Replication Appliance



Adds a new replication appliance. This option is only valid for agentless vSphere protection.

View Server Details



Views detailed information about a server. This button leaves the **Manage Servers** page and opens the **View Server Details** page. See *Viewing server details* on page 288

Remove Server



Removes the server from the console.

Provide Credentials



Changes the login credentials that the Double-Take Console use to authenticate to a server. This button opens the **Provide Credentials** dialog box where you can specify the new account information. See *Providing server credentials* on page 287. You will remain on the **Manage Servers** page after updating the server credentials.

Manage Group Assignments



Allows you to assign, move, and remove the selected server from specific server groups. This buttons opens the Manage Group Assignments dialog box where you can assign and unassign the server to specific server groups. The server will appear in server groups marked with a checkmark, and will not appear in groups without a checkmark. Servers assigned to a server group will automatically appear in parent server groups.



Installs or upgrades Double-Take on the selected server. This button opens the Install page where you can specify installation options. See Installing using the Double-Take Console on page 231.

Uninstall 🎉

Uninstalls Double-Take on the selected server.

View Server Events



The option is not available for Linux sources or appliances.

View Server Logs



Views the Double-Take logs messages for a server. This button opens the Logs window. This separate window allows you to continue working in the Double-Take Console while monitoring log messages. You can open multiple logging windows for multiple servers. When the Double-Take Console is closed, all logging windows will automatically close.

Refresh

Refreshes the status of the selected servers.

Gather Support Diagnostics

Executes the diagnostic DTInfo utility which collects configuration data for use when reporting problems to technical support. It gathers Double-Take log files; Double-Take and system settings; network configuration information such as IP, WINS, and DNS addresses; and other data which may be necessary for technical support to troubleshoot issues. You will be prompted for a location to save the resulting file which is created with the information gathered. Because this utility is gathering several pieces of information, across the network to your console machine, it may take several minutes to complete the information gathering and sending the resulting file to the console machine.

View Replication Service Details

Views the replication service details for a server. This option is not applicable to Linux source servers or appliances.

Overflow Chevron



Displays any toolbar buttons that are hidden from view when the window size is reduced.

Providing server credentials

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

Viewing server details

The **View Server Details** page allows you to view details about that particular server. The server details vary depending on the type of server or appliance you are viewing.

Server name

The name or IP address of the server. If you have specified a reserved IP address, it will be displayed in parenthesis.

Operating system

The server's operating system version

Roles

The role of this server in your Double-Take environment. In some cases, a server can have more than one role.

- EngineRole—Source or target server
- ProxyRole—A Linux appliance for a full server to ESX appliance job
- ProxiedRole—A Linux source server for a full server to ESX appliance job
- ControllerRole—Controller appliance for an agentless vSphere job
- ReplicationApplianceRole—Replication appliance for an agentless vSphere job
- Reporting Service—Double-Take Reporting Service server

Status

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

Activity

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

Connected via

The IP address and port the server is using for communcations. You will also see the Double-Take protocol being used to communicate with server. The protocol will be XML web services protocol (for servers running Double-Take version 5.2 or later) or Legacy protocol (for servers running version 5.1 or earlier).

Version

The product version information

Access

The security level granted to the specified user

User name

The user account used to access the server

Licensing

Licensing information for the server

- Warning or error icon—Warnings indicate the license is temporary and will
 expire. Errors indicate the license has expired or it is invalid for the operating system
 on the server.
- **Product**—The product associated with the license
- Serial Number—The serial number associated with the license
- Expiration Date—The date the license expires, if there is one
- Activation Code—The activation code
- **Licensing Status**—The status of the license on the server. If your license is expired, any jobs using that server will be in an error state.
- License Type—The type of Double-Take license

Editing server properties

Highlight a server on the **Manage Servers** page and click **View Server Details** from the toolbar. Under **Tasks**, select **Edit server properties**. The **Edit Server Properties** page allows you to view and edit properties for that server. Click on a heading on the **Edit Server Properties** page to expand or collapse a section of properties.

• Server licensing on page 291—Views, adds, and removes activation codes

Server licensing

Licensing identifies your Double-Take activation codes.



The fields and buttons in the **Licensing** section will vary depending on your Double-Take Console configuration and the type of activation codes you are using.

Add activation codes and activation keys—The activation code and activation key are the
Double-Take license which is required on every Double-Take server. They are a 24 character,
alpha-numeric code. You can change your activation code without reinstalling, if your license
changes.

To add an activation code and activation key, type in the code and click **Add**. If your console has been enabled to manage your license inventory, click **Choose from inventory** to open the Activation Codes dialog box where you can select the activation codes you want to apply. See *Console options* on page 277 for details on enabling the license inventory.



The license inventory feature cannot be enabled if your service provider has restricted access to it.

- Current activation codes—The server's current activation codes are displayed.
 - **Warning or error icon**—Warnings indicate the license is temporary and will expire. Errors indicate the license has expired.
 - **Product**—The product associated with the license
 - Serial Number—The serial number associated with the license
 - Expiration Date—The date the license expires, if there is one
 - Activation Code—The activation code

To remove a code, highlight it and click **Remove**. To copy a code, highlight it and click **Copy**.

- Activation—If your activation code needs to be activated, you will see an additional Activation section at the bottom of the Licensing section. To activate your code, use one of the following procedures.
 - Obtain activation key online, then activate—If you have Internet access, click the
 hyperlink in the Activation section to take you to the web so that you can submit your
 activation information. Complete and submit the activation form, and you will receive an email with the activation key. Activate your server by entering the activation key in the Add
 activation codes and activations keys field and clicking Add.
 - Obtain activation key offline, then activate—If you do not have Internet access, go to https://activate.doubletake.com from another machine that has Internet access. Complete and submit the activation form, and you will receive an e-mail with the activation key. Activate your server by entering the activation key in the Add activation codes and activations keys field and clicking Add.

The permanent code is specific to this server. It cannot be used on any other server. If the activation code and server do not match, Double-Take will not run.



If your activation codes needs to be activated, you will have 14 days to do so.

If you rename a server that has already has a Double-Take license applied to it, for example if you rebuild a server, you will have to perform a host-transfer to continue using that license. This includes changing the case (capitalization) of the server name (upper or lower case or any combination of case).

Viewing server logs

You can view the Double-Take and Double-Take Management Service log files through the Double-Take Console using either of these two methods.

- On the Manage Servers page, highlight a server in the list and click View Server Logs from the toolbar.
- On the Manage Jobs page, right-click a job and select View Logs. Select either the source server log or the target server log.

Separate logging windows allow you to continue working in the Double-Take Console while monitoring log messages. You can open multiple logging windows for multiple servers. When the Double-Take Console is closed, all logging windows will automatically close.

The following table identifies the controls and the table columns in the **Server logs** window.



This button starts the addition and scrolling of new messages in the window.



This button pauses the addition and scrolling of new messages in the window. This is only for the **Server logs** window. The messages are still logged to their respective files on the server.



This button copies the messages selected in the **Server logs** window to the Windows clipboard.

Clear 2

This button clears the **Server logs** window. The messages are not cleared from the respective files on the server. If you want to view all of the messages again, close and reopen the **Server logs** window.

Filter

From the drop-down list, you can select to view all log messages or only those messages from the Double-Take engine or the Double-Take Management Service.

Time

This column in the table indicates the date and time when the message was logged.

Description

This column in the table displays the actual message that was logged.

Service

This column in the table indicates if the message is from the Double-Take engine or the Double-Take Management Service.

License management and activation

- License management—You can manage your Double-Take licenses through the license
 inventory feature in the Double-Take Console. Ideally, you should select one machine in your
 organization where you want to maintain the license inventory because the license inventory does
 not communicate between machines.
 - From the license inventory, you can add and remove Double-Take licenses manually. You can import and export a Double-Take license file to handle groups of licenses. You can also activate the licenses in your inventory. See *Managing the Double-Take license inventory* on page 296 for details on using the license inventory.
- **Licensing a server or appliance**—If you did not license your server during the installation (appliances are not licensed during the installation), you can apply a Double-Take license to a server or appliance using the console. See *Licensing a server or appliance* on page 298.
- Activation—For those Double-Take licenses that require activation, you have 14 days to activate
 the license. You can use the license inventory to activate multiple licenses at one time, or you can
 activate your licenses one at a time on a per server basis.

Managing the Double-Take license inventory

You can manage your Double-Take licenses through the license inventory feature in the Double-Take Console. Ideally, you should select one machine in your organization where you want to maintain the license inventory because the license inventory does not communicate between machines.

By default, the license inventory feature is enabled. You can, and should, disable it on machines that are not maintaining your license inventory. To disable the license inventory, select **Options** from the console toolbar, deselect **Enable license inventory**, and click **Save**.

To manage your license inventory, select **Go**, **Manage License Inventory**.



The license inventory feature may not appear in your console if your service provider has restricted access to it.

You will see the following fields on the **Manage License Inventory** page.

Warning or error icon

Warnings indicate the license has not been activated, or it is temporary and will expire. Errors indicate the license has expired.

Serial Number

The serial number associated with the license

Product

The Double-Take product associated with the license

License Type

A short description of the type of license, including any expiration date or quantity information

Version

The product version number associated with the license

Server

The name of the server the license has been applied to, if any. If the same activation code is used on multiple servers, for example with evaluation licenses, you will see multiple entries in the license inventory for each server that is using that activation code.

Use the following toolbar controls to manage your licenses.

Import Licenses



Imports all of the activation codes from a license inventory file into the license inventory. This is a file you may have received from Vision Solutions, or it may be from another Double-Take Console. Depending on the information contained in the file, you may be prompted to activate some of the servers in your license inventory.

Export Licenses



Add Licenses



Allows you to manually enter your activation codes. Enter the activation codes in the Add Licenses dialog box in the space provided, separating multiple codes by a comma or by putting each code on a separate line. Once your activation codes are in your license inventory, you can apply them to a server in several ways. See Activating a single license on page 300 or See Activating multiple licenses on page 302.

Remove License



Removes the selected activation code from the license inventory. You can only remove activation codes that are not being used by any server in the console.

Reclaim License



Returns a **Single** license type back into the console's license inventory. You may want to reclaim a single license if you have removed a license from a server so that you can assign it to another server. This process should not be used for transferring an **Activated** license to a new or rebuilt server, which is considered deactivation.

Activate



Activates the activation codes in the license inventory. For complete details on this process, see Activating multiple licenses on page 302.

Licensing a server or appliance

If you did not license your server during the installation (appliances are not licensed during the installation), you can apply a Double-Take license to a server or appliance using the console.

- 1. Make sure you have your server or appliance inserted in the console.
- 2. From the Manage Servers page, double-click on the server to view the server's details.
- 3. From the View Server Details page, click on the Edit server properties link.
- 4. Expand the **Licensing** section.
- 5. Licensing identifies your Double-Take activation codes.



The fields and buttons in the **Licensing** section will vary depending on your Double-Take Console configuration and the type of activation codes you are using.

Add activation codes and activation keys—The activation code and activation key are
the Double-Take license which is required on every Double-Take server. They are a 24
character, alpha-numeric code. You can change your activation code without reinstalling, if
your license changes.

To add an activation code and activation key, type in the code and click **Add**. If your console has been enabled to manage your license inventory, click **Choose from inventory** to open the Activation Codes dialog box where you can select the activation codes you want to apply. See *Console options* on page 277 for details on enabling the license inventory.



The license inventory feature cannot be enabled if your service provider has restricted access to it.

- Current activation codes—The server's current activation codes are displayed.
 - **Warning or error icon**—Warnings indicate the license is temporary and will expire. Errors indicate the license has expired.
 - **Product**—The product associated with the license
 - Serial Number—The serial number associated with the license
 - Expiration Date—The date the license expires, if there is one
 - Activation Code—The activation code

To remove a code, highlight it and click **Remove**. To copy a code, highlight it and click **Copy**.

- Activation—If your activation code needs to be activated, you will see an additional
 Activation section at the bottom of the Licensing section. To activate your code, use one
 of the following procedures.
 - Obtain activation key online, then activate—If you have Internet access, click
 the hyperlink in the Activation section to take you to the web so that you can submit
 your activation information. Complete and submit the activation form, and you will

receive an e-mail with the activation key. Activate your server by entering the activation key in the **Add activation codes and activations keys** field and clicking **Add**.

 Obtain activation key offline, then activate—If you do not have Internet access, go to https://activate.doubletake.com from another machine that has Internet access. Complete and submit the activation form, and you will receive an e-mail with the activation key. Activate your server by entering the activation key in the Add activation codes and activations keys field and clicking Add.

The permanent code is specific to this server. It cannot be used on any other server. If the activation code and server do not match, Double-Take will not run.



If your activation codes needs to be activated, you will have 14 days to do so.

If you rename a server that has already has a Double-Take license applied to it, for example if you rebuild a server, you will have to perform a host-transfer to continue using that license. This includes changing the case (capitalization) of the server name (upper or lower case or any combination of case).

6. Once you have completed your licensing, click **OK** to return to the **Manage Servers** page.

Activating a single license

You will need to have Double-Take installed on a single server, then you can activate the license on that server using the Double-Take Console. You can repeat this process for other servers individually.

- 1. Click Manage Servers in the toolbar.
- 2. Click Add Servers in the Manage Servers page toolbar.
- Specify the name of your server and the credentials of the account used to install Double-Take, and click Add.
- 4. Click OK.
- 5. After you server has been added to the **Manage Servers** page, click **View Server Details** in the toolbar.
- 6. On the View Server Details page, click Edit server properties under Tasks.
- Expand the **Licensing** section.



- 8. If your activation code needs to be activated, you will see an additional **Activation** section at the bottom of the **Licensing** section. To activate your code, use one of the following procedures.
 - Obtain activation key online, then activate—If you have Internet access, click the
 hyperlink in the Activation section to take you to the web so that you can submit your

- activation information. Complete and submit the activation form, and you will receive an e-mail with the activation key. Activate your server by entering the activation key in the **Add** activation codes and activations keys field and clicking **Add**.
- Obtain activation key offline, then activate—If you do not have Internet access, go to https://activate.doubletake.com from another machine that has Internet access. Complete and submit the activation form, and you will receive an e-mail with the activation key. Activate your server by entering the activation key in the Add activation codes and activations keys field and clicking Add.

The permanent code is specific to this server. It cannot be used on any other server. If the activation code and server do not match, Double-Take will not run.



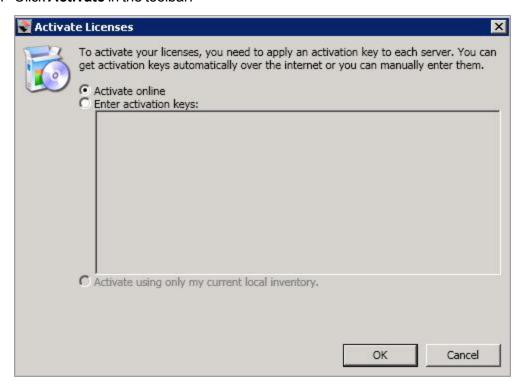
If your activation codes needs to be activated, you will have 14 days to do so.

If you rename a server that has already has a Double-Take license applied to it, for example if you rebuild a server, you will have to perform a host-transfer to continue using that license. This includes changing the case (capitalization) of the server name (upper or lower case or any combination of case).

Activating multiple licenses

You will need to have Double-Take installed on multiple servers, then you can activate the licenses on those servers all at once.

- 1. On the **Manage Servers** page, verify that all of the Double-Take servers that you want to activate have been added to your Double-Take Console.
- 2. Select Go, Manage License Inventory.
- 3. Click Activate in the toolbar.



- Enter activation keys—Select this option on a non-Internet connected machine to
 manually enter activation keys. You can obtain these keys from
 https://activate.doubletake.com by entering the server information for each server manually
 or by uploading an export file of your license inventory. See Activating a single license on
 page 300 for details on gathering the server information or Managing the Double-Take
 license inventory on page 296 for details on creating an export file.
- Activate using only my current local inventory—Select this option to activate all of the
 licenses in your inventory based on the inventory's current local settings. For example, if a
 server was offline when you performed activate online, you can perform the activation
 process again when the server is available. You may also need to perform this option if you
 selected not to activate the servers when you imported a license inventory file.
- 4. Click **OK** to begin the activation process.

After your servers have been activated, you will have permanent codes that are specific to each server. They cannot be used on any other server. If the activated code and server do not match, Double-Take will not run.



If your activation codes needs to be activated, you will have 14 days to do so.

If you rename a server that has already has a Double-Take license applied to it, for example if you rebuild a server, you will have to perform a host-transfer to continue using that license. This includes changing the case (capitalization) of the server name (upper or lower case or any combination of case).

Managing VMware servers

The Manage VMware Server page allows you to view, add, remove, or edit credentials for your VMware servers available in the console.

VMware Server

The name of the VMware server

Full Name

The full name of the VMware server

User Name

The user account being used to access the VMware server

Add VMware Server



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

Remove Server



Remove the VMware server from the console.

Provide Credentials



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

Index	console 231, 275-276, 296 credentials 287
	Linux source 233 options 277 servers 283, 288, 290
.rpm 26	virtual recovery appliance 225 core operations 8
A	create job full server to ESX appliance 236
activation 295, 298, 300, 302 activation code 43, 190 activation codes 291, 298	credentials 287 CustomerCare 2
adding Linux source 233 adding virtual recovery appliance 225	D
appliance license 228, 298 auto-disconnect 119 auto-reconnect 119, 121 auto-remirror 56, 127	daemon 43 data failover monitoring 175 monitoring 64 data protection 50 database files 203
В	disconnect connection 123 documentation 43
bandwidth limiting 162 block device adding 53, 136, 141 configuration 43	dtfs_mounts 43 dtloop 43 DTSetup 43 DTStat 266
С	E
chained configuration 16 compression 164 configurations 11 connection 51, 56	e-mail notification 92, 206 erasing 29
database storage file 203	F
disconnect 123 firewall 60 ID 113 NAT 60 overview 113 reconnect 121 simulation 62 target processing 122 types 50 Connection Manager 56 Connection Wizard 51	failback overview 166 failover 10 configuring failover monitoring 167 editing failover monitoring configuration 174 failover 178 full server to ESX appliance 274 monitoring 175 overview 166 removing failover monitoring configuration 178

Failover Control Center 44, 48 refresh rate 49 file differences mirror 56, 127 options compared 125 file system 20 configuration 43 firewall 60 full server to ESX appliance create job 236 edit job 261 failover 274	kernel type 20 kernel version 20 L legal 2 license 228, 291, 296, 298 licensing 190 Linux requirements 214
job details 257 job log 263 job validation 260 managing 250 requirements 214 source 237	logging 279, 293 e-mailing system messages 92, 206 log file 70, 74, 204 messages 75 verification 154 logging on and off Replication Console 210
G	·
gettting started 213	M
H hard links 132	managing jobs full server to ESX appliance 250 managing servers 283, 304 manual intervention 167, 178 many to one configuration 14
ICMP 60 identity 188 installation console 223, 231 Linux 229 Linux servers 26 overview 24 virtual recovery appliance 219 Windows client 28, 223	mirroring 8, 56 automatically 127 controls 125 overview 124 monitoring data workload 64 failover data 175 monitoring tools overview 63 SNMP 105
J	N
job full server to ESX appliance 236, 250, 257, 260-261, 274	named pipes 210 NAT 60 network communications 195
job log 263	0
	one to many configuration 15

one to one configuration 12 orphan files 129 overview 7-8, 11, 213, 275	requirements 20 console 276 full server to ESX appliance 214 Linux 214
P	resources 2
ports 43, 60, 195 protection 50 full server to ESX appliance 236 overview 7, 18	S security credentials 47
Q	groups 43 overview 209 server details 288
queues 119 overview 114 queuing data 116, 197	server identity 188 server properties 290-291 server settings 187 single server configuration 17 SNMP
reconnecting automatically 121 recurse subdirectories 134, 139 remirror 56 removing 29 replication 9 capabilities 132 configuration 43 inserting tasks 149 overview 131 starting 148 Replication Console 44-45 credentials 47 logging on and off 210 workspaces 46 replication set block device 141	configuration 106 overview 105 statistics 110 traps 107 soft links 132 source 7 data processing options 200 full server to ESX appliance 237 server startup options 193 statistics 265-266, 268 file 96-97, 204 output 99 overview 95 SNMP 110 synchronization 8 system messages e-mail 92, 206
calculating size 144 copying 143 creating 53, 136 creating manually 139 database storage file 203 deleting 147 limitiations 134 modifying 142 overview 134 renaming 143 rules 139	target 7 data procesing options 202 pause 122 resume 122 task command processing 149 TDU 62 technical support 2

```
transmission
bandwidth limiting 162
controls 157
network communications 195
overview 156
schedule database storage file 203
scheduling 157
```

U

UDP 60
uninstallation 29
upgrade
console 231
Linux 229
Linux servers 26
overview 24
virtual recovery appliance 219

٧

verification
log file 154, 204
manually verifying 151
overview 150
scheduling 152
virtual systems 20

W

workspaces 46