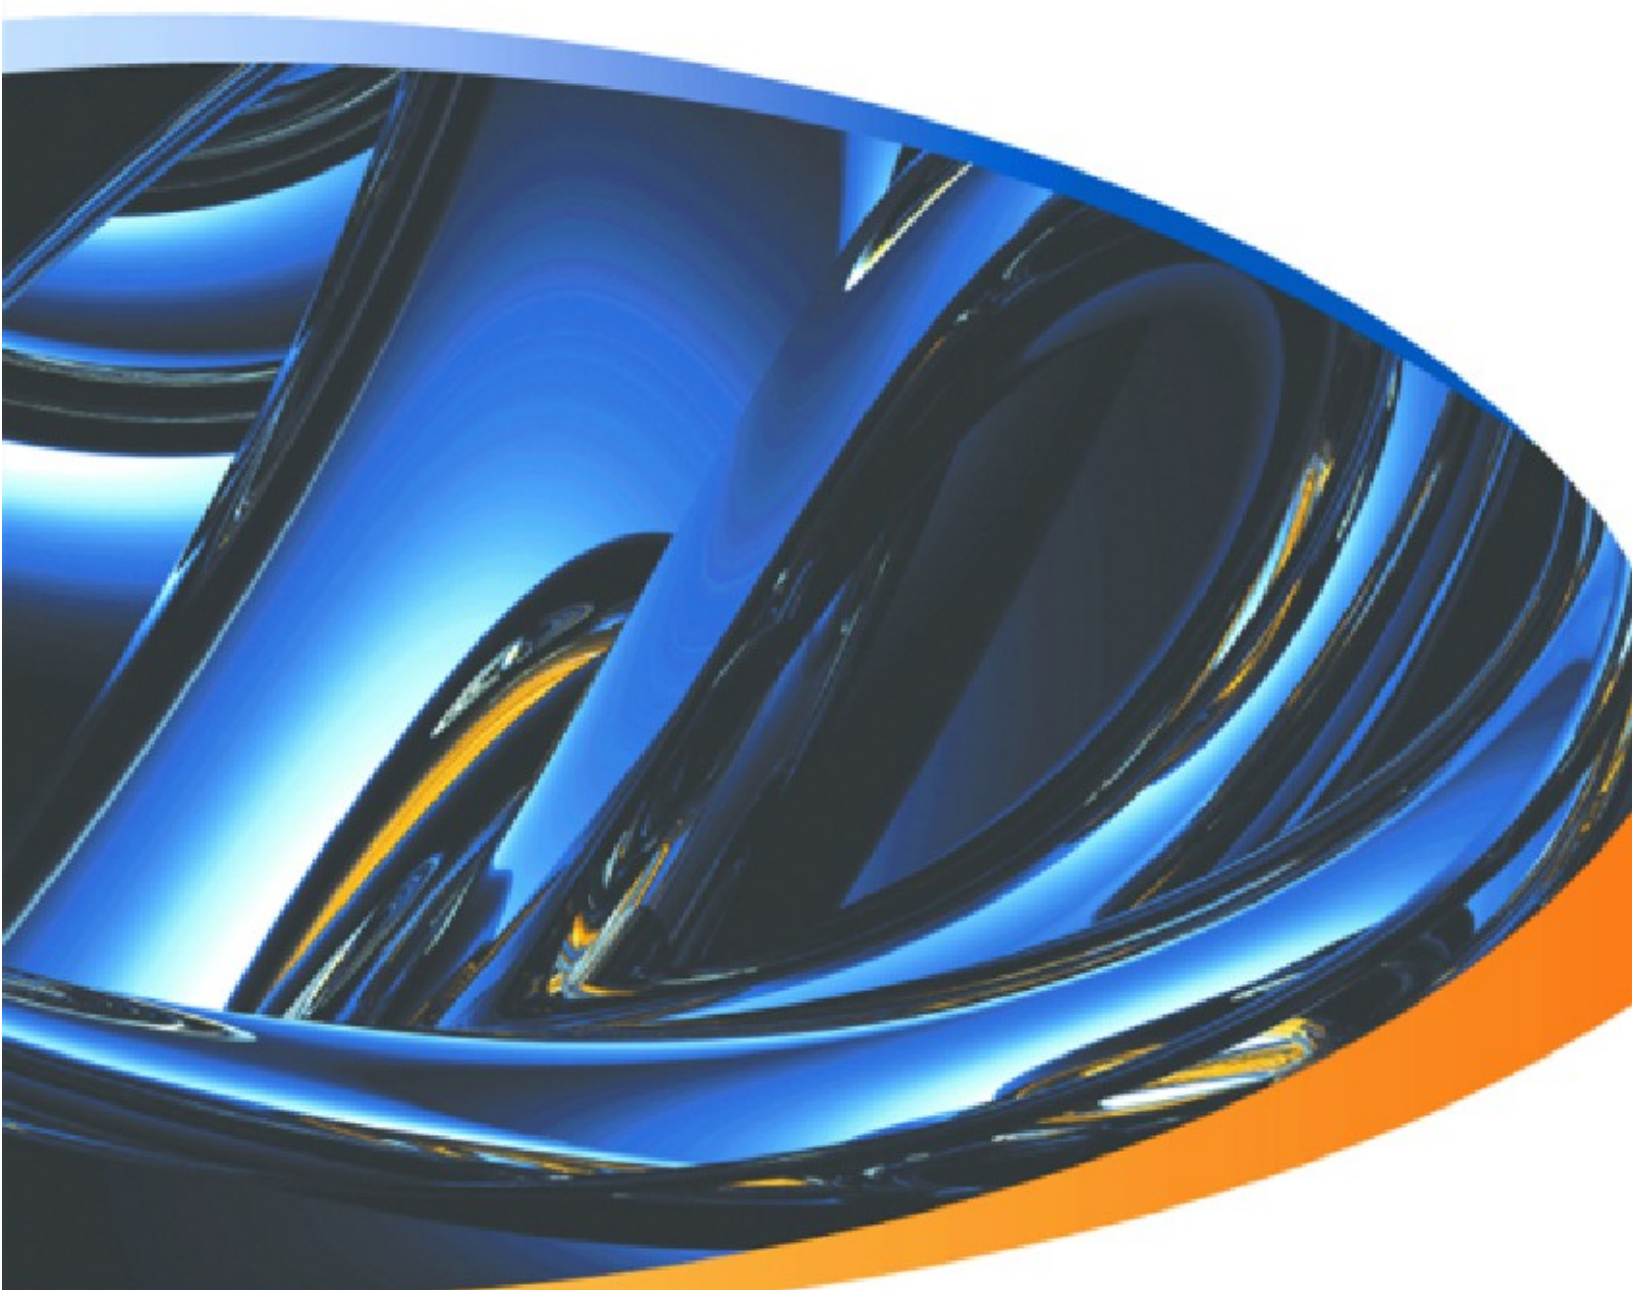


Double-Take Backup User's Guide



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

Table of Contents

Welcome to the Double-Take Backup User's Guide	1
Getting started.....	1
Frequently used pages.....	1
Resources.....	1
Double-Take Backup overview	2
TimeData continuous data protection.....	3
Ontrack PowerControls.....	3
Cargo archiving.....	4
Double-Take Backup requirements	5
Source server requirements.....	5
Repository server requirements.....	6
Recovery server requirements.....	7
TimeData requirements.....	8
Ontrack PowerControls requirements.....	8
Cargo requirements.....	9
Console requirements.....	10
Replication capabilities.....	10
Installation	13
Installation and upgrade notes.....	13
Installing or upgrading Double-Take Backup.....	14
Installing Double-Take Backup automatically.....	16
Installing TimeData continuous data protection.....	19
Upgrading to SQL Server 2008 Standard.....	20
Installing and configuring Ontrack PowerControls.....	20
Managing server licenses.....	21
Double-Take Backup console	23
Importing a job.....	23
Clearing console security credentials.....	24
Source protection	25
Protecting your source.....	25
Optional protection settings.....	29
Compressing data for transmission for a protection job.....	30
Limiting bandwidth for a protection job.....	30
Routing transmission for a protection job.....	30
E-mailing repository server event messages.....	31
Data on the repository server	33
Viewing and retrieving data from snapshots or from a TimeData point in time.....	33
Viewing and retrieving Exchange data.....	35
Recovery	36
Installing Double-Take Backup on the recovery server.....	36
Recovering an entire server to an existing physical or virtual machine.....	37

Recovering an entire server to an automatically provisioned virtual machine on VMware ESX.....	41
Recovering an entire server to an automatically provisioned virtual machine on Hyper-V.....	48
Recovering data to an existing physical or virtual machine.....	53
Optional recovery settings.....	56
Compressing data for transmission for a recovery job.....	57
Limiting bandwidth for a recovery job.....	57
Routing transmission for a recovery job.....	57
E-mailing recovery server event messages.....	58
Job monitoring.....	60
Monitoring a job through the Double-Take Backup console.....	60
Monitoring log files.....	67
Monitoring event messages.....	68
Event messages.....	68
E-mailing source event messages.....	99
Statistics.....	100
Viewing the statistics file.....	101
Statistics.....	102
Performance Monitor.....	108
Monitoring Performance Monitor statistics.....	108
Performance Monitor statistics.....	108
SNMP.....	111
Configuring SNMP on your server.....	112
SNMP traps.....	112
SNMP statistics.....	114
Error codes.....	117
Cargo archiving.....	123
Configuring archiving security.....	123
Running an archive preview report.....	124
Manually archiving data.....	126
Archiving files individually.....	126
Archiving groups of files based on criteria.....	126
Archiving data on a schedule.....	127
Choosing a storage location for archived files on the repository server.....	128
Recalling archived files.....	129
Recalling files individually.....	129
Recalling all files.....	130
Using firewalls.....	131
Double-Take Backup ports.....	131
Microsoft Windows ports.....	132
Hardware ports.....	132
Index.....	133

Welcome to the Double-Take Backup User's Guide

This online documentation is designed to help you find detailed information about Double-Take Backup. Use the **Contents** pane on the left to explore the different sections or use **Search** or **Index** to find specific topics. Below are some of the most useful and most accessed topics in the online documentation.

If you are using Internet Explorer and are getting an active content blocker pop-up in your browser, you need to enable active content to run in files on your computer. From Internet Explorer, select **Tools, Internet Options**. On the **Advanced** tab, enable **Allow active content to run in files on My Computer**.

Getting started

- [Double-Take Backup overview](#) explains the basic functionality of Double-Take Backup.
- [Double-Take Backup requirements](#) lists the minimum requirements for each Double-Take Backup server, the Double-Take Backup console, and the TimeData and Cargo features.
- [Installation](#) includes installation instructions.

Frequently used pages

- [Source protection](#) includes step-by-step instructions for configuring source protection.
- [Recovery](#) contains links to the various recovery methods.
- [Job monitoring](#) details the different methods for monitoring a protection or recovery job.

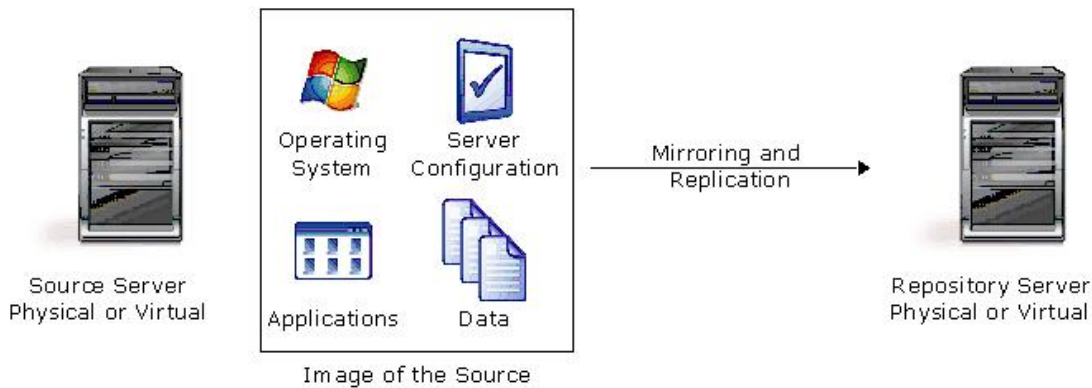
Resources

- **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 [Double-Take Software support web site](#).
- **Sales**—If you need maintenance renewal, an upgrade activation code, or other sales assistance, contact your reseller/distributor or a Double-Take Software sales representative. Contact information is available on the [Double-Take Software web site contact information page](#).
- **Technical Support**—If you need technical assistance, you can contact technical support. All basic configurations outlined in the online documentation will be supported through technical support. Your technical support center is dependent on the reseller or distributor you purchased your product from and is identified on your service agreement. If you do not have access to this agreement, contact Double-Take Software Technical Support and they will direct you to the correct service provider. To contact Double-Take Software Technical Support, you will need your serial number and activation code. Contact information is available on the [Double-Take Software web site support information page](#).
- **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 [Double-Take Software Professional Services information page](#).
- **Training**—Classroom and computer-based training are available. For more information, see the [Double-Take Training information page](#).
- **Documentation**—Please forward any comments or suggestions about this online documentation to documentation@doubletake.com.

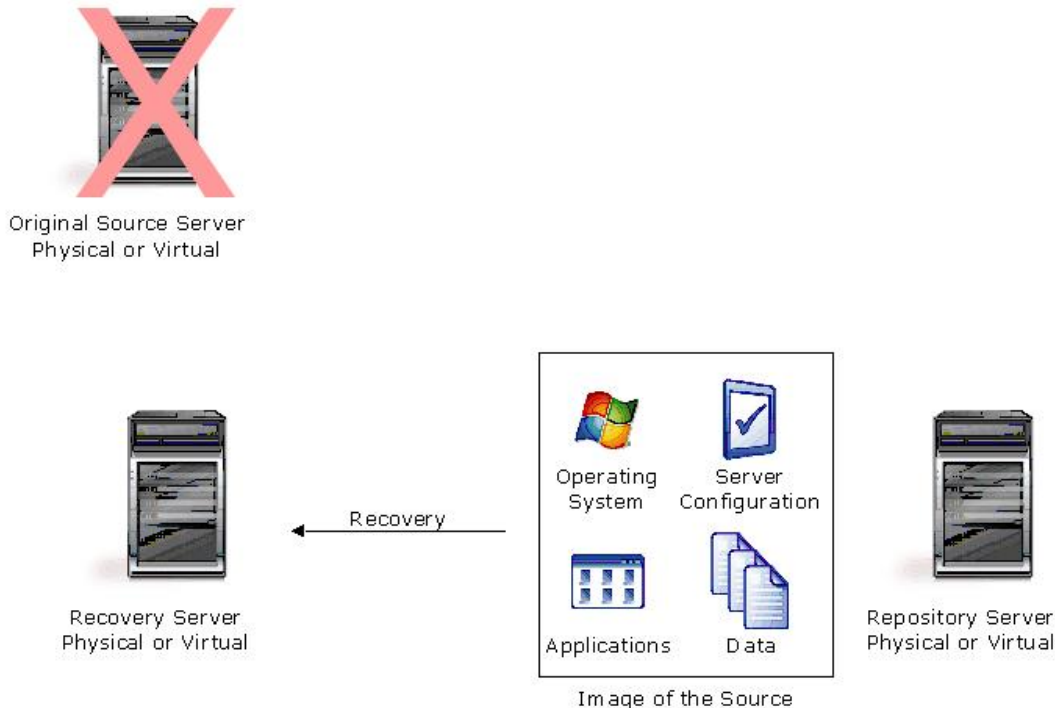
Double-Take Backup overview

Double-Take Backup is a comprehensive disk-to-disk backup and recovery solution. It allows you to protect a production server, known as a source, by replicating data from the source or an image of the entire source to another server, known as a repository server. The repository server can be located locally or over a WAN. The image of an entire source contains all of the source server's data and the source server's system state, which is the server's configured operating system and the installed applications. The data or entire server image can be recovered to a dissimilar physical or virtual environment.

As changes are made on the source, Double-Take Backup mirroring and replication keeps the image of the source stored on the repository server up-to-date. Double-Take Backup replicates, in real-time, only the file changes rather than copying entire files, allowing you to more efficiently use resources.



In the event the source should fail, you can use the image of the source, stored on the repository server, to quickly and easily create a new source, which can be a physical or virtual machine.




Double-Take Backup also provides features that enhance or extend the benefits of disk-to-disk protection.

- [TimeData continuous data protection](#) enhances the recoverability of data by providing any-point-in-time recovery of data or entire servers.
- [Ontrack PowerControls](#), an optional feature, adds Exchange data recovery allowing you to recover messages, folders, attachments, mailboxes, and so on to an Exchange Server or a .pst file.
- [Cargo source disk optimization](#) saves disk space on the source server by archiving old files to the repository server.
- Atom file-level deduplication reduces storage requirements on the repository server by storing duplicate files only once. This feature will be available in a future release.

TimeData continuous data protection

Double-Take Backup provides the TimeData continuous data protection feature that records every change made to the source image. With every change recorded on the repository server, recovery can occur from any point within the continuous data protection window specified by the user.

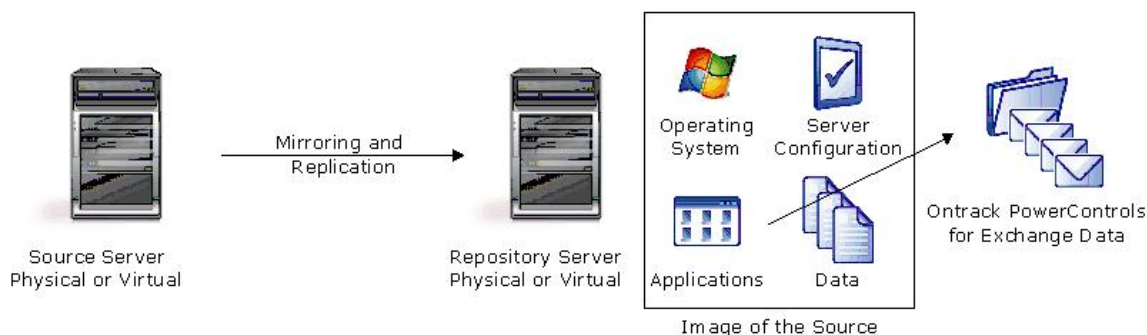
TimeData is installed and runs on the repository server. It keeps track of all changes to the source image. Changes are temporarily stored in the Event cache until they can be processed and stored in the Event log.



The longer TimeData stores all changes, the more disk space is consumed. To reduce disk space usage you should keep changes for only as long as needed and rely on Double-Take Backup snapshots for older point-in-time-recovery. For example, retain the last seven days of all changes using TimeData continuous data protection and rely on Double-Take Backup snapshots for older points in time.

Ontrack PowerControls


If you purchased a Double-Take Backup Repository with Ontrack PowerControls license, you can recover Exchange data including messages, folders, attachments, mailboxes, and so on to an Exchange Server or a .pst file.



Ontrack PowerControls can recover Exchange data from the repository server from live data, a Double-Take Backup snapshot, or from a TimeData continuous data protection window.

Cargo archiving

Cargo reduces the amount of disk space used on a source server by identifying files on the source that have not been accessed recently and converting those files to links to the repository server. The conversion process deletes the contents of the file from the source server, creates a reparse point to the full contents of the file on the repository server, and marks the file as sparse. In simpler terms, the source file becomes an empty pointer to the full contents of the file located on the repository server. When an archived file on the source is accessed, the full contents are recalled seamlessly from the repository server with no interruption to the requesting user or application.



Double-Take Backup can archive files based on the age of files, allowing you to keep newer files fully on the source. You can exclude files from archiving based on size, which is useful for excluding small files where the archiving benefit would be minimal.

When a file has been archived, the link file is still accessible on the source. The link file has an offline file attribute associated with it. When an end-user attempts to access the link file, the contents of the file are recalled from the repository server back to the source. The file opens to the end-user as if it had always been fully on the source.

To help you plan archiving and fine-tune your archive settings, you can generate a preview report. The archiving preview report does not actually archive any data but reports how much data would be archived based on the settings you select.

Double-Take Backup requirements

Each Double-Take Backup server must meet minimum requirements. Additionally, the machine where you will be running the Double-Take Backup console must also meet some basic requirements.

- **Source server**—This is the server that contains the data you will be protecting. Verify your server meets the [source server requirements](#).
- **Repository server**—This is the server that will be protecting your source server(s). Verify your server meets the [repository server requirements](#).
 - **TimeData** —The TimeData feature can be run on your repository server and has additional minimum requirements for the repository server. TimeData records every change made to the source image(s) on the repository server allowing you to recover from any point within the TimeData window. Verify your server meets the [repository server](#) and [TimeData requirements](#).
 - **Ontrack PowerControls**—The Ontrack PowerControls feature can be run on your repository server and has additional minimum requirements for the repository server. Ontrack PowerControls allows you to recover Microsoft Exchange messages, folders, attachments, mailboxes, and so on to an Exchange server or a .pst file. Verify your server meets the [repository server](#) and [Ontrack PowerControls requirements](#).
- **Recovery server**—This is the server that will become your new source when you are recovering an entire server. This server can be an existing physical or virtual machine or Double-Take Backup can provision (automatically create) a virtual machine for you during the recovery process. Verify your server meets the [recovery server requirements](#).
- **Cargo** —Cargo can be run on your source(s) and repository server. Cargo archives old files on the source to the repository server, thus reducing the amount of disk space used on the source. Verify your servers meets the [source server](#), [repository server](#), and [Cargo requirements](#).
- **Console**—The Double-Take Backup console is used to manage your Double-Take Backup servers. Verify the machine where you are running the console meets the [console requirements](#).

In addition to the server and feature requirements, review the [Replication capabilities](#) to understand the type of data that Double-Take Backup protects.

Source server requirements

This is the server that contains the data you will be protecting.

- **Operating system**—Double-Take Backup can protect a physical or virtual source server running any of the following operating system editions.
 - Windows Server 2008 or 2008 R2 Datacenter, Enterprise, Standard, Essential Business Server, Web Server, Foundation Server, Small Business Server, or Storage Server Edition
 - Windows Server 2003 or 2003 R2 Datacenter, Enterprise, Standard, Web Edition, Small Business Server, or Storage Server Edition. Each of the Windows 2003 operating systems require Service Pack 1 or later.
- **Architecture**—32-bit and 64-bit architectures are supported
- **File system**—Double-Take Backup supports the same file system formats that Microsoft supports: FAT, FAT32, and NTFS.
- **System memory**—There are different memory requirements depending on the system architecture you are using. Be sure you have at least the minimum amount of memory for your environment. You should consider having at least the recommended amount of system memory.

Architecture	Minimum System Memory	Recommended System Memory
32-bit	128 MB	At least 512 MB
64-bit	512 MB	At least 1024 MB

- **Disk space for program files**—This is the amount of disk space, approximately 130 MB, needed on the source to hold the Double-Take Backup program files.
- **Server name**—Double-Take Backup includes Unicode file system support, but your server name must be in ASCII format.
- **Network protocols**—Your servers must use TCP/IP with static IP addressing.
- **Microsoft .NET Framework**—Version 3.5 Service Pack 1 is required on the source.
- **Junction points**—Only domain controller junction points will be replicated. Any user-created junction points will be replicated as two unlinked copies of the data.

Repository server requirements

This is the server that will be protecting your source server(s).

- **Operating system**—The repository server can be a physical or virtual server running any of the following operating system editions.

Note: If you are protecting an entire server, your repository server must have a same or newer operating system than the source (not including service pack level). For example, you cannot protect a Windows 2008 source to a Windows 2003 repository server. This is because the repository server must be able to process the source's registry for system state processing.

- Windows Server 2008 or 2008 R2 Datacenter, Enterprise, Standard, Essential Business Server, Web Server, Foundation Server, Small Business Server, or Storage Server Edition
- Windows Server 2003 or 2003 R2 Datacenter, Enterprise, Standard, Web Edition, Small Business Server, or Storage Server Edition. Each of the Windows 2003 operating systems require Service Pack 1 or later.
- **Architecture**—32-bit and 64-bit architectures are supported
- **File system**—Double-Take Backup supports the same file system formats that Microsoft supports: FAT, FAT32, and NTFS.
- **System memory**—There are different memory requirements depending on the system architecture you are using. Be sure you have at least the minimum amount of memory for your environment. You should consider having at least the recommended amount of system memory, especially if your repository server is protecting multiple source servers.

Architecture	Minimum System Memory	Recommended System Memory
32-bit	128 MB	At least 1024 MB
64-bit	512 MB	At least 2048 MB

- **Disk space for program files**—This is the amount of disk space, approximately 130 MB, needed on the repository server to hold the Double-Take Backup program files.
- **Disk space for data files**—This is the amount of disk space on the repository server to hold all of the source data files. This will be dependent on the number of sources you are protecting, the applications you are running on each source, and the amount of data files on each source.
- **Disk space for system state image**—This is the amount of disk space on the repository server to hold the image of the source system state. You should have at least 5 GB of space on the repository server for the system state. The more applications you are running on the source, the larger the system state image will be. Therefore, you may need additional space if you source has many applications. If your repository server is protecting multiple sources, you will need space for each source's system state.

- **Server name**—Double-Take Backup includes Unicode file system support, but your server name must be in ASCII format.
- **Mount point limitations**—You cannot store your source data or source image on the repository server on a volume that is accessed by a mount point.
- **Network protocols**—Your servers must use TCP/IP with static IP addressing.
- **Microsoft .NET Framework**—Version 3.5 Service Pack 1 is required on the repository server.

Recovery server requirements

This is the server that will become your new source when you are recovering an entire server. This server can be an existing physical or virtual machine or Double-Take Backup can provision (automatically create) a virtual machine for you during the recovery process.

- **Operating system**—You have several options for your recovery server operating system.
 - **Physical or existing virtual**—If your recovery server is a physical server or an existing virtual server, it can have any of the following Windows operating system editions. A physical or existing virtual recovery server must have the same operating system as the original source.
 - Windows Server 2008 or 2008 R2 Enterprise, Standard, Web, Small Business Server, Enterprise x64, or Standard x64
 - Windows Server 2003 or 2003 R2 Enterprise, Standard, Web, Small Business Server, Enterprise x64, or Standard x64. Each of the Windows 2003 operating systems require Service Pack 1 or later.
 - **Provisioned virtual**—If your recovery server will be provisioned (automatically created during the recovery process) you must have a physical host machine where Double-Take Backup can create the new virtual server. Your physical host machine can be any of the following operating systems.
 - Windows 2008 Hyper-V
 - The Enterprise Edition of VMware ESX version 3.0.x, 3.5, or 4.0. If you are using version 3.0.1, you must have HotRemove patch #6921838. VMware ESXi versions 3 or 4 are also supported as long as ESXi is managed by VCenter or VSphere.
- You must have an existing virtual machine, known as a virtual recovery appliance, running Windows Server 2008, 2008 R2, 2003, or 2003 R2. The virtual recovery appliance must have a same or newer operating system than the original source (not including service pack level).
- The existing virtual must have Double-Take Backup installed and licensed on it before you can start the recovery process.

When you are ready for recovery the existing virtual will create the new virtual server, mount disks, format disks, and so on. Once the new virtual machine is online, it will have the identity, data, and system state of the original source. Since the existing virtual machine maintains its own identity, it can be reused for additional recoveries.
- **System memory**—There are different memory requirements depending on the system architecture you are using. Be sure you have at least the minimum amount of memory for your environment. You should consider having at least the recommended amount of system memory.

Architecture	Minimum System Memory	Recommended System Memory
32-bit	128 MB	At least 512 MB
64-bit	512 MB	At least 1024 MB

- **Disk space for program files**—This is the amount of disk space, approximately 130 MB, needed on the recovery server to hold the Double-Take Backup program files.

- **Disk space for data files**—This is the amount of disk space on the recovery server to hold all of the source data files. This will be dependent on the applications you are running and the amount of data files you have.
- **Disk space for system state image**—This is the amount of disk space on the recovery server to hold the image of the source system state. You should have at least 5 GB of space on the recovery server for the system state. The more applications you are running on the source, the larger the system state image will be. Therefore, you may need additional space if your source has many applications.
- **Disk types**—Thin-provisioned virtual disks on VMware ESX are not supported because they cannot be shared by two virtual machines simultaneously.
- **Server name**—Double-Take Backup includes Unicode file system support, but your server name must be in ASCII format.
- **Network protocols**—Your servers must use TCP/IP with static IP addressing.
- **Microsoft .NET Framework**—Version 3.5 Service Pack 1 is required on the recovery server.

TimeData requirements

In addition to the [repository server requirements](#), the following requirements apply to the TimeData continuous data protection functionality of Double-Take Backup.

- **Operating system**—If your repository server is running a Windows 2003 operating system, you must have Service Pack 2 installed on the server.
- **File system**—NTFS is the only supported file system. Do not use FATx file systems on the repository server.
- **System memory**—The minimum system memory requirements for the repository server are higher with TimeData continuous data protection.

Architecture	Minimum System Memory	Recommended System Memory
32-bit	1 GB	2 GB
64-bit	2 GB	4 GB

- **Disk space for program files**—This is the amount of disk space, approximately 1.5 GB, needed on the repository server to hold the TimeData program files.
- **Disk space for TimeData historical data**—This is the amount of disk space, approximately two and a half times the amount of data being protected for each source, needed on the repository server to hold the TimeData continuous data protection historical data. You may need to adjust the amount of disk space needed based on the applications and data on your source, the number of days of historical data you choose to store, and the change rate of your data. There is a combined limit of 2 TB of TimeData historical data for all sources that you are protecting to the repository server.
- **Microsoft Internet Information Services**—Microsoft IIS version 5.0 or later is required on the repository server.
- **Microsoft ASP.NET**—The version of Microsoft ASP.NET included with your Windows operating system is required on the repository server.
- **Microsoft SQL Server**—TimeData will automatically install Microsoft SQL Server Express or Standard, depending on your Double-Take Backup media (CD or web download).
- **Drive configuration**—For best results using TimeData continuous data protection, you should separate the operating system, SQL Server, the TimeData drive, Event log, and Event cache on different drives.

Ontrack PowerControls requirements

Ontrack PowerControls supports recovery from the following versions of Exchange.

-
- Exchange Server 5.5
 - Exchange Server 2000
 - Exchange Server 2003
 - Exchange Server 2007

Ontrack PowerControls can recover from an unlimited number of Exchange servers.

In addition to the [repository server requirements](#), the storage files (.edb files) that Ontrack PowerControls will open must be accessible locally or over the network.

Cargo requirements

In addition to the [source](#) and [repository server requirements](#), the following requirements apply to the Cargo archiving functionality of Double-Take Backup.

- **Operating system**—The following operating system and file system requirements apply to archiving.
 - You must use the NTFS file system. FAT and FAT32 are not supported for archiving.
 - Transactional files on Windows 2008 are not supported. Any request for a recall within a transaction will generate a sharing violation error and the recall will not occur. The application can open and read a byte outside of a transaction to cause the recall to occur and then proceed with the transaction.
 - Do not modify the system attributes on any file or folder. Incorrect system attributes could cause data to be skipped for archiving, or it could cause source crashes because the full system files must remain available on the source.
 - Because the archiving feature uses the file access date and time, do not disable the Windows system setting that updates file access date and time. If this setting is disabled, the archiving functionality will not behave as expected.
 - Archiving can only be performed on non-system volumes. (The system volume is where the Windows system files are located.)
- **Network configuration**—The following network configuration requirements apply to archiving.
 - Workgroup environments are not supported. The source and repository server should be member servers in a domain environment.
 - The source can be a domain controller as long as Active Directory and system data is not selected for archiving.
- **Security**—Specific [archiving security configuration](#) must be completed before using the archiving functionality.
- **Anti-virus**—You must configure your anti-virus software to skip archived files or the archived files will be recalled during virus scans. Additionally, you must configure your anti-virus software so that the Double-Take Backup application is a low-risk process. Configure the low-risk processes so that they do not scan when writing to or reading from the disk. If you need assistance configuring your anti-virus software, see your virus software documentation. Keep in mind that your anti-virus software may refer to archived files differently, for example, offline, migrated, or sparse files.
- **Cargo functionality**—Review the following requirements, limitations, and informational notes before using the archiving feature.
 - Archiving supports a one-to-one or many-to-one source to repository server configuration. A one-to-many protection job is not supported for archiving because if there were multiple links, the source would not know which repository server to recall from.
 - Double-Take Backup will not archive the following file types.
 - System files
 - Alternate data streams
 - Encrypted files

-
- Reparse points
 - Archiving is not a continuous process. It can be initiated manually or a schedule can be established.
 - Only data that is being protected by Double-Take Backup will be archived. If you exclude data when configuring your source protection job, that excluded data will not be eligible for archiving.
 - Modifying the attributes of a link file on the source will not recall the contents of the file from the repository server. The attribute changes will be replicated to the copy of the file on the repository server, but the source will remain a link file.
 - Double-Take Backup automatically creates a DTArchiveBin share on the repository server for archiving activity. Do not modify the name or security permissions of the share. If Double-Take Backup does not find the expected share name and permissions, users will be unable to recall their data.
 - The DTArchiveBin maintains the only copy of your archived files. You may want to consider regularly backing up the copy of the data on the repository server and the DTArchiveBin, for example to a tape backup.
 - When an archived file is deleted, the link on the source is deleted. The archived content in the DTArchiveBin will be deleted also if there is an active protection job present with replication enabled.
 - Briefcase files cannot be archived and will cause an error in the Double-Take Backup log. These files should be added to the list of files excluded from the archive process.
 - Cargo archiving is not compatible with [TimeData](#) continuous data protection. If you use one, do not use the other.

Console requirements

The Double-Take Backup console is used to manage your Double-Take Backup servers.

- **Operating system**—The Double-Take Backup console can be run from a source or repository server. It can also be run from a Windows XP or Vista machine.
- **Microsoft .NET Framework**—Version 3.5 Service Pack 1 is required to run the console.

Replication capabilities

Double-Take Backup replicates file and directory data stored on any Windows file system (FAT, FAT32, NTFS4, and NTFS5). Replicated items also include Macintosh files, compressed files, NTFS attributes and ACLs (access control list), dynamic volumes, files with alternate data streams, sparse files, and encrypted files. Files can be replicated across mount points, even though mount points are not created on the repository server. Some reparse points are replicated, including CommVault Data Migrator and BridgeHead Software HT FileStore.

Typically, Double-Take Backup does not replicate items that are not stored on the file system, such as physical volume data and registry based data. Nor does it replicate NTFS extended attributes, registry hive files, Windows or any system or driver pagefile, system metadata files (\$LogFile, \$Mft, \$Bitmap, \$Extend\\\$UsnJrnl, \$Extend\\\$Quota, \$Extend\\\$ObjId, and \$Extend\\\$Reparse), hard links, or the Double-Take Backup disk-based queue logs. However, if you are protecting a full-server, Double-Take Backup will automatically gather and replicate all necessary system state data, including files for the operating system and applications.

Note the following replication caveats.

1. If you have mixed file systems, keep in the mind the following.
 - a. If, on your source, you have a FAT volume mounted on a directory which resides on an NTFS volume, these files will not be mirrored, regardless of the target file system. Replication will work correctly. To work around this issue, make sure both volumes are NTFS.

-
- b. If you are mirroring/replicating from an NTFS source to a FAT target, you may see additional error messages in your Double-Take Backup log file because the target file system cannot handle the NTFS attributes or file permissions. For example, if your protection job contains files with alternate data streams, you will see messages indicating that there are unfinished operations because the FAT file system cannot store the alternate data stream information.
 - c. If you select a compressed file or folder from an NTFS partition and replicate it to a FAT target, the attributes are lost, but the data is maintained.
 2. If any directory or file contained in your protection job that specifically denies permission to the system account or the account running the Double-Take service, the attributes of the file on the target will not be updated because of the lack of access. This also includes denying permission to the Everyone group because this group contains the system account.
 3. If you select a dynamic volume and you increase the size of the volume, the target must be able to compensate for an increase in the size of the dynamic volume.
 4. If you select files with alternate data streams, keep in mind the following.
 - a. Alternate data streams are not included in the protection job size calculation. Therefore, you may see the mirror process at 100% complete while mirroring continues.
 - b. The number of files and directories reported to be mirrored will be incorrect. It will be off by the number of alternate streams contained in the files and directories because the alternate streams are not counted. This is a reporting issue only. The streams will be mirrored correctly.
 - c. Use the checksum option when performing a difference mirror or verification to ensure that all alternate data streams are compared correctly.
 - d. If your alternate streams are read-only, the times may be flagged as different if you are creating a verification report only. Initiating a remirror with the verification will correct this issue.
 5. If you select encrypted files, keep in mind the following.
 - a. Only the data, not the attributes or security/ownership, is replicated. However, the encryption key is included. This means that only the person who created the encrypted file on the source will have access to it on the target.
 - b. Only data changes cause replication to occur; changing security/ownership or attributes does not.
 - c. Replication will not occur until the Windows Cache Manager has released the file. This may take awhile, but replication will occur when Double-Take Backup can access the file.
 - d. When remirroring, the entire file is transmitted every time, regardless of the remirror settings.
 - e. Empty encrypted files will be mirrored to the target, but if you copy or create an empty encrypted file within the protection job after mirroring is complete, the empty file will not be created on the target. As data is added to the empty file on the source, it will then be replicated to the target.
 - f. When you are replicating encrypted files, a temporary file is created on both the source and target servers. The temporary file is automatically created in the same directory as the Double-Take Backup disk queues. If there is not enough room to create the temporary file, an out of disk space message will be logged. This message may be misleading and indicate that the drive where the encrypted file is located is out of space, when it actually may be the location where the temporary file is trying to be created that is out of disk space.
 6. If you are using mount points, keep in mind the following.
 - a. By default, the mount point data will be stored in a directory on the target. You can create a mount point on the target to store the data or maintain the replicated data in a directory. If you use a directory, it must be able to handle the amount of data contained in the mount point.
 - b. Recursive mount points are not supported. If you select data stored on a recursive mount point, mirroring will never finish.
 7. Double-Take Backup supports transactional NTFS (TxF) write operations, with the exception of TxF SavePoints (intermediate rollback points).
-

-
- a. With transactional NTFS and Double-Take Backup mirroring, data that is in a pending transaction is in what is called a transacted view. If the pending transaction is committed, it is written to disk. If the pending transaction is aborted (rolled back), it is not written to disk.

During a Double-Take Backup mirror, the transacted view of the data on the source is used. This means the data on the target will be the same as the transacted view of the data on the source. If there are pending transactions, Double-Take Backup will indicate this state. As the pending transactions are committed or aborted, Double-Take Backup mirrors any necessary changes to the target. Once all pending transactions are completed, the state will be updated.

If you see the pending transactions state, you can check the Double-Take Backup log file for a list of files with pending transactions. As transactions are committed or aborted, the list is updated until all transactions are complete.
 - b. During replication, transactional operations will be processed on the target identically as they are on the source. If a transaction is committed on the source, it will be committed on the target. If a transaction is aborted on the source, it will be aborted on the target.
 - c. When recovery occurs any pending transactions on the target will be aborted before the source identity is assigned to the target.
8. Double-Take Backup supports Windows 2008 symbolic links and junction points. A symbolic link is a link (pointer) to a file. Junction points are also links, but to folders and volumes.
 - a. If the link and the file/folder/volume are both in your source protection job, both the link and the file/folder/volume are mirrored and replicated to the target.
 - b. If the link is in the source protection job, but the file/folder/volume it points to is not, only the link is mirrored and replicated to the target. The file/folder/volume that the link points to is not mirrored or replicated to the target. A message is logged to the Double-Take Backup log identifying this situation.
 - c. If the file/folder/volume is in the source protection job, but the link pointing to it is not, only the file/folder/volume is mirrored and replicated to the target. The link pointing to the file/folder/volume is not mirrored or replicated to the target.
 9. Short file names are not supported on FAT file systems.

Installation

Review the [Double-Take Backup requirements](#) and [Installation and upgrade notes](#) before beginning your installation. You have the following installation options available.

- [Installing or upgrading Double-Take Backup](#)—Use these instructions to install or upgrade Double-Take Backup using the installation wizard.
- [Installing Double-Take Backup automatically](#)—Use these instructions to install or upgrade Double-Take Backup using command-line parameters which allow you to automate the process by running an unattended installation.
- [Installing TimeData continuous data protection](#)—If you are installing on a repository server and will be using TimeData continuous data protection, you will need to install TimeData after you install Double-Take Backup. If you entered a repository server activation code during the Double-Take Backup installation, the TimeData installation can be launched directly from the Double-Take Backup installation. Otherwise, you will need to start the TimeData installation manually.
- [Upgrading to SQL 2008 Standard](#)—The TimeData installation includes SQL Server. If you installed from a web download, the SQL Server version is 2008 Express edition. The SQL database is limited to 4 GB in the Express edition. If you installed from the Double-Take Backup CD, the SQL Server version is 2008 Standard edition. The SQL database is limited to 524272 TB in the Standard edition. If you have the Double-Take Backup CD, you can upgrade a 2008 Express edition to the 2008 Standard edition.
- [Installing and configuring Ontrack PowerControls](#)—If you purchased a Double-Take Backup Repository with Ontrack PowerControls license, you can install Ontrack PowerControls on your repository server. If you entered an appropriate Ontrack PowerControls activation code during the Double-Take Backup installation, the Ontrack PowerControls installation can be launched directly from the Double-Take Backup installation. Otherwise, you will need to start the Ontrack PowerControls installation manually.

Installation and upgrade notes

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

- Since Double-Take Backup installs device drivers, it is recommended that you update your Windows Recovery Disk, before installing or making changes to your servers. For detailed instructions on creating a recovery disk, see your Windows reference manuals. Make sure that you select the option to back up the registry when building the repair disks.
- Because Double-Take Backup 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 Backup installed, you must remove Double-Take Backup prior to the operating system upgrade. Uninstall Double-Take Backup, perform the operating system upgrade, and then reinstall Double-Take Backup.
- If you are installing to a drive other than the drive which contains your system TEMP directory, the Microsoft Windows Installer will still load approximately 100 MB of data to the TEMP directory during the installation. If you do not have enough disk space on the drive that contains the TEMP directory, you may need to change where it is located.
- During installation, a file called dtinfo.exe is installed to the Double-Take Backup installation directory. This program can be run to collect configuration data for use when reporting problems to technical support. It gathers Double-Take Backup log files; Double-Take Backup and system registry settings; network configuration information such as IP, WINS, and DNS addresses; and other data which may be necessary for customer support to troubleshoot issues. After running the executable, a zip file is automatically created with the information gathered.
- Double-Take Backup 5.2 can interoperate back to version 5.0 but has the following limitations. The Double-Take Backup clients can only control the same or earlier releases. To accommodate rolling upgrades, older sources can connect to newer targets, but newer sources cannot connect to older targets. Configurations not listed in the following chart are not supported.

Client	Source	Target	Supported
5.0	5.1 or 5.2	5.1 or 5.2	No
5.1	5.2	5.2	No
5.2	5.0	5.0, 5.1, or 5.2	Yes
5.2	5.1	5.0	No
5.2	5.1	5.1 or 5.2	Yes
5.2	5.2	5.0 or 5.1	No
5.2	5.2	5.2	Yes

- When performing a rolling upgrade, update the target server first. After the upgrade is complete, any sources will automatically reconnect to the target. Upgrade the source(s) when convenient.

Installing or upgrading Double-Take Backup

Use these instructions to install Double-Take Backup or upgrading an existing Double-Take Backup installation.

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

-
6. Review and accept the Double-Take Software license agreement to continue with the installation program. Click **Next** to continue.
 7. Select the type of installation you would like to perform on this machine.
 - **Client and Server Components**—This options installs both the client and server components. The server components are required for systems that will function as a source, repository server, or recovery server. The server requires an activation code for the service to run. The client does not require an activation code, but it is required to administer this and other servers throughout the organization.
 - **Client Components Only**—This option installs only the client components. The client components do not require an activation code, but are required to administer servers throughout the organization.
 - **Server Components Only**—This option installs only the server components. The server components are required for systems that will function as a source, repository server, or recovery server. The server requires an activation code for the service to run.
 8. If desired, specify where the Double-Take Backup files will be installed.
 9. Click **Next** to continue.

Note: If you selected a **Client Components Only** installation type, continue with step 16 below.

10. You will be prompted to enter your activation code information. Your activation code is a 24-character, alpha-numeric activation code which applies the appropriate license to your installation. Multiple codes allow you to run multiple Double-Take Software products.
 - To add a new code to this machine, enter the code and click **Add**. Repeat this process if you have additional codes.
 - To remove any codes from this machine, highlight the code in the **Current activation codes** list and click **Remove**.
11. When you have added your activation code(s), click **Next** to continue.
12. Double-Take Backup uses system memory to store data in queues. Specify the maximum amount of system memory to be used for the Double-Take Backup queues and click **Next** to continue.

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

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

-
18. Depending on the activation code(s) you entered during the installation and if you selected a **Client and Server Components** or **Server Components Only**, you may see one or two additional installation options after the Double-Take Backup files have completed copying. The installation options are for TimeData continuous data protection and Ontrack PowerControls for Exchange messages, folders, and mailboxes. If you want to install either of these features, enable the options and the [TimeData](#) and [Ontrack PowerControls](#) installations will automatically start when the Double-Take Backup installation closes. If both options are selected, the Ontrack PowerControls installation will run first.
 19. Click **Finish** to exit the installation program.

Installing Double-Take Backup automatically

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

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

Sample DTSetup.ini

```
[Config]
DTSETUPTYPE=DTNT
DTACTIVATIONCODE=123456789012345678901234
DOUBLETAKEFOLDER="C:\Program Files\Double-Take Software\Double-Take"
QMEMORYBUFFERMAX=128
DISKQUEUEFOLDER="C:\Program Files\Double-Take Software\Double-Take"
DISKQUEUEMAXSIZE=UNLIMITED
DISKFREESPACEMIN=50
DTSERVICESTARTUP=1
PORT=6320
SET_FWPORT=Y
```

```
[Alpha]
DTSETUPTYPE=DTNT
DTACTIVATIONCODE=123456789012345678901234
DOUBLETAKEFOLDER="C:\Program Files\Double-Take Software\Double-Take"
QMEMORYBUFFERMAX=128
DISKQUEUEFOLDER="C:\Program Files\Double-Take Software\Double-Take"
DISKQUEUEMAXSIZE=UNLIMITED
DISKFREESPACEMIN=50
DTSERVICESTARTUP=1
PORT=6320
SET_FWPORT=Y
```

```
[Beta]
DTSETUPTYPE=DTNT
DTACTIVATIONCODE=123456789012345678901234
```

```

DOUBLETAKEFOLDER="C:\Program Files\Double-Take Software\Double-Take"
QMEMORYBUFFERMAX=128
DISKQUEUEFOLDER="C:\Program Files\Double-Take Software\Double-Take"
DISKQUEUEMAXSIZE=UNLIMITED
DISKFREESPACEMIN=50
DTSERVICESTARTUP=1
PORT=6320
SET_FWPORT=Y

```

In the sample DTSetup file, the server Alpha would use the parameters defined under the [Alpha] heading. The server Beta would use the parameters defined under the [Beta] heading. All other servers would use the configuration under the [Config] section.

Note: Do not leave any parameter blank in the Config or MachineName sections. Otherwise, a hard-coded default stored in a Double-Take Backup .dll file will be used.

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

Parameter	Valid Values
DTSetupType	<ul style="list-style-type: none"> • DTNT—Both the Double-Take Backup server and client components will be installed. The server components are required for systems that will function as a source, repository server, or recovery server. The server requires an activation code for the service to run. The client does not require an activation code, but it is required to administer this and other servers throughout the organization. • DTCO—Only the Double-Take Backup client components will be installed. The client components do not require an activation code, but are required to administer servers throughout the organization. • DTSO—Only the Double-Take Backup server components will be installed. The server components are required for systems that will function as a source, repository server, or recovery server. The server requires an activation code for the service to run.
DTActivationCode	A 24 character, alpha-numeric activation code which applies the appropriate license to the server
DoubleTakeFolder	Any valid path specifying the location of the Double-Take Backup files
QMemoryBufferMax	Any integer representing the amount of system memory, in MB, to use for memory-based queuing
DiskQueueFolder	Any valid path to the location of the disk-based queue.
DiskQueueMaxSize	Any integer representing the amount of disk space, in MB, to use for disk-based queuing or the keyword UNLIMITED which will allow the queue usage to automatically expand whenever the available disk space expands
DiskFreeSpaceMin	Any integer representing the amount of disk space, in MB, that must remain free at all times

Parameter	Valid Values
DTSERVICESTARTUP	<ul style="list-style-type: none"> • Y or 1—Start the Double-Take service automatically • N or 0—Do not start the Double-Take service automatically This parameter is not applied if your DTSetupType is DTCO.
Port	Any integer between 1024 and 65535 that identifies the Windows Firewall port used for Double-Take Backup.
Set_FWPORT	<ul style="list-style-type: none"> • Y or 1—Set the Windows Firewall port for Double-Take Backup • N or 0—Do not set the Windows Firewall port for Double-Take Backup

You must have Microsoft .NET installed or enabled on the server before starting the automatic installation.

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

If you are installing to a repository server and will be using TimeData continuous data protection, the automatic installation will not automatically launch the TimeData installation. See [Installing TimeData continuous data protection](#) for instructions.

Installing or upgrading automatically to a local machine

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

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

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

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

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

Installing or upgrading automatically to a remote machine

1. Create a temporary directory on the primary site server. For example, create z:\dtinstall.
2. Share the temporary folder.

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

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

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

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

Note: Substitute your mapped drive for m:\.

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

```
C:\>net use m: \\server_name\share
```

```
The command completed successfully
```

```
C:\>M:
```

```
M:\>setup_1352 /s /v"DTSETUPINI=\"m:\DTSetup.ini\" /qn"
```

Installing TimeData continuous data protection

1. Depending on your Double-Take Backup installation selections, the TimeData installation may automatically launch immediately following the Double-Take Backup installation. If it did not start automatically or if you are installing it separately from the Double-Take Backup installation, start the TimeData installation program using the appropriate instructions, depending on your media source.
 - **CD**—Load the Double-Take Backup CD into the local CD-ROM drive. If auto-run is enabled, the installation program will start automatically. To manually start the program, select **Start, Run** and specify <cd_drive>:\setup\TimeData\TimeData.exe.
 - **Web download**—Unzip the .exe file that you downloaded from the web to a temporary location and select **Start, Run** and specify <temporary_unzip_location>\setup\TimeData\TimeData.exe.
2. At the Welcome screen, click **Next**.
3. Review and accept the license agreement to continue with the installation program. Click **Next** to continue.
4. Specify your **User Name** and **Organization** and click **Next** to continue.
5. Select the directory where you want to install TimeData and click **Next** to continue.

For best results, you should separate the various TimeData components that you will be prompted to install (TimeData Event log, TimeData drive letter, SQL Server, and TimeData Event cache) on different drives.

-
6. Specify the location to store the TimeData Event log. For best results, do not select a drive containing Windows system files.
 7. Select the drive letter to indicate where the TimeData continuous data protection historical files will be stored. A TS_Drive network share will be created on the drive. The drive will be accessible through Windows Explorer and other applications.
 8. Click **Next** to continue.
 9. Select the directory where you want to install the SQL Server databases used with TimeData and click **Next** to continue.
 10. Specify the location to store the TimeData Event cache, which temporarily stores data until it can be stored in the TimeData Event log.
 11. Specify the maximum size of the TimeData Event cache. By default, the maximum size is 20 GB. Make sure you have sufficient disk space in the selected location for the Event cache.
 12. Click **Next** to continue.
 13. At the Ready screen, click **Install**.
 14. During the installation, you may be prompted to add an exception to the Windows Firewall for TimeData. Click **OK** to add the port exception. If you **Cancel** the port modification, you will have to manually modify your firewall settings for TimeData processing.
 15. Click **Finish** to exit the installation program. If prompted, reboot your computer for the installation changes to complete.

Upgrading to SQL Server 2008 Standard

The TimeData installation includes SQL Server. If you installed from a web download, the SQL Server version is 2008 Express edition. The SQL database is limited to 4 GB in the Express edition. If you installed from the Double-Take Backup CD, the SQL Server version is 2008 Standard edition. The SQL database is limited to 524272 TB in the Standard edition. If you have the Double-Take Backup CD, you can upgrade a 2008 Express edition to the 2008 Standard edition.

1. Load the Double-Take Backup CD into the local CD-ROM drive. If auto-run is enabled, the Setup Launcher will start automatically. Close the Setup Launcher when it opens.
2. Locate and run the file \setup\TimeData\SQLServer2008Upgrade.exe on the CD.
3. At the Welcome screen, click **Next**.
4. Review and accept the license agreement to continue with the installation program. Click **Next** to continue.
5. Specify your **User Name** and **Organization** and click **Next** to continue.
6. At the Ready screen, click **Install**.
7. When the upgrade has completed, click **Finish** to exit the installation program.

Installing and configuring Ontrack PowerControls

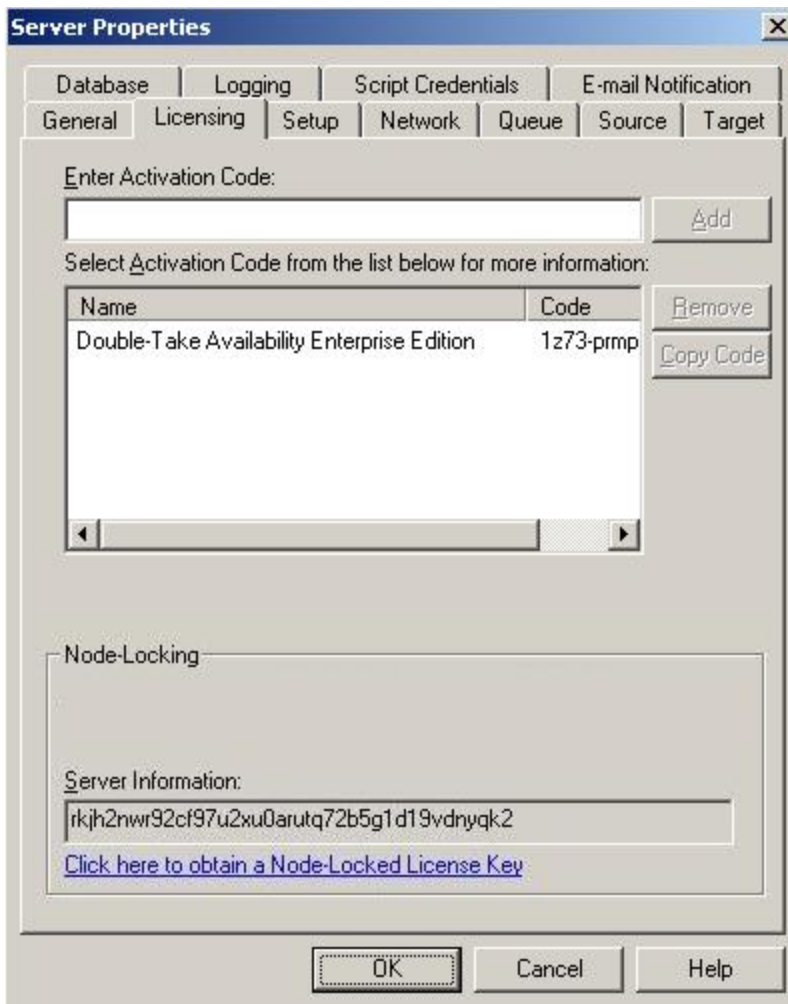
1. Depending on your Double-Take Backup installation selections, the Ontrack PowerControls installation may automatically launch immediately following the TimeData installation. If it did not start automatically or if you are installing it separately from the Double-Take Backup installation, start the Ontrack PowerControls installation program using the appropriate instructions, depending on your media source.
 - **CD**—Load the Double-Take Backup CD into the local CD-ROM drive. If auto-run is enabled, the installation program will start automatically. To manually start the program, select **Start, Run** and specify <cd_drive>:\setup\Ontrack PowerControls\Ontrack PowerControls.exe.
 - **Web download**—Unzip the .exe file that you downloaded from the web to a temporary location and select **Start, Run** and specify <temporary_unzip_location>\setup\Ontrack PowerControls\Ontrack PowerControls.exe.
2. At the Welcome screen, click **Next**.
3. Review and accept the license agreement by clicking **Yes**.
4. Specify your **User Name** and **Company Name** and click **Next** to continue.

-
5. Select the directory where you want to install Ontrack PowerControls and click **Next** to begin the installation.
 6. Click **Finish** to exit the installation program.
 7. After the installation is complete, you need to configure Ontrack PowerControls to use your repository server. Select **Start, Programs, Kroll Ontrack, Ontrack PowerControls for Double-Take, Ontrack PowerControls for Exchange**.
 8. Cancel the Data Wizard when it appears.
 9. Select **Help, About**, and click **License Info**.
 10. Specify the name of your Double-Take Backup repository server. You can also click the browse button (the button with ... on it) to select the name from a network drill-down list.
 11. Click **Close** to save the server information.
 12. Click **OK** to close the About box.

Managing server licenses

From the Replication Console, you can manage your server activation codes. The activation code is the Double-Take Backup license which is required on every Double-Take Backup 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 Backup from being used illegally on multiple servers, you may have received a node-locked activation code, which is a temporary license. The temporary license is not activated until you login to the server. Once the temporary license is activated, you have 14 days to update it to a permanent, node-locked license. The permanent node-locked license will be created by supplying unique server information to Double-Take Software. Since the permanent node-locked license contains unique server information, specific to the hardware where Double-Take Backup is installed, the node-locked license cannot be used on any other server, thus prohibiting illegal applications.
1. Open the Replication Console by selecting **Start, Programs, Double-Take, Double-Take Replication Console**.
 2. Double-click your Double-Take Backup server in the left pane of the Replication Console to log on to it.
 3. Right-click the server name and select **Properties**.
 4. Select the **Licensing** tab. The fields displayed on this tab will vary depending on your activation code.

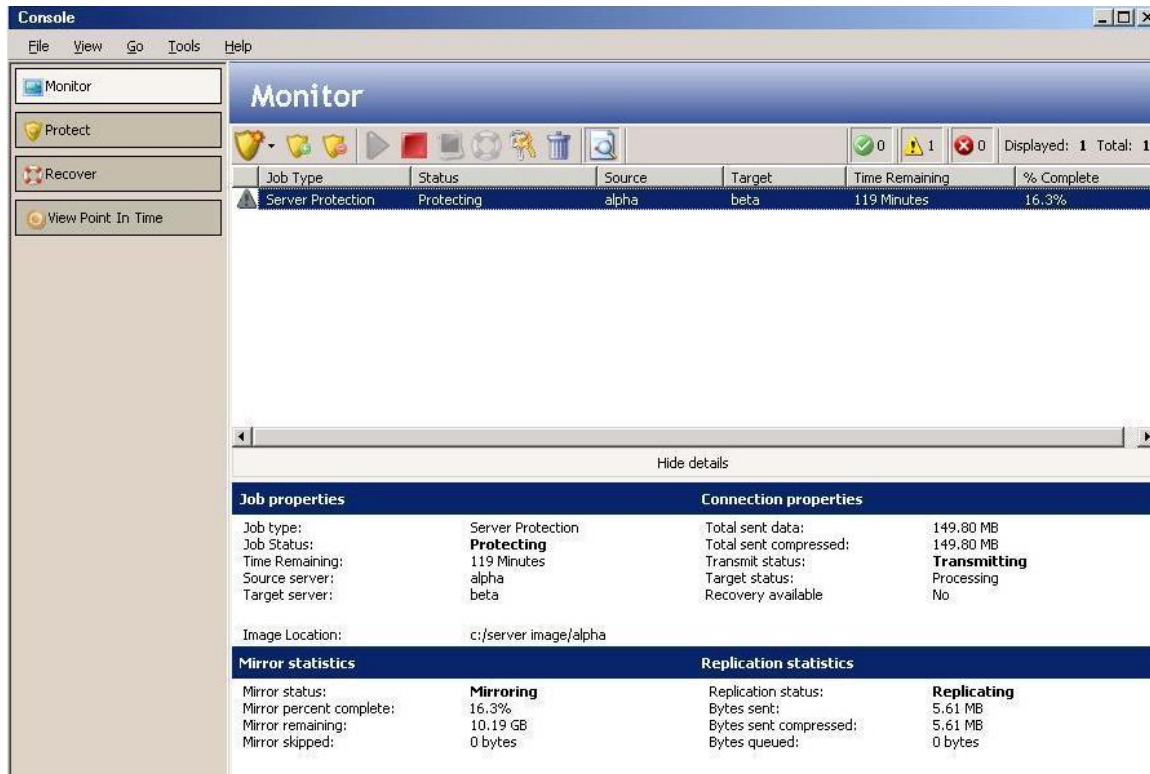


5. Enter an activation code and click **Add**. Repeat for each activation code.
6. If you need to remove a code from the server, highlight it in the list and click **Remove**.
7. To update a temporary node-locked license to a permanent license, you need to provide server information which will be used to generate a permanent node-locked license.
 - a. Highlight your activation code in the list to display the Node-Locking section at the bottom of the **Licensing** tab.
 - b. 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.
 - c. 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 Backup will not run.
8. Click **OK** to save the settings.

Double-Take Backup console

The Double-Take Backup console is the user interface that manages, monitors, and controls your backup solutions. To open the console, select **Start, Programs, Double-Take, Double-Take Backup, Double-Take Backup Console**.

On the left pane of the console are static navigation buttons which initiate workflows and display monitoring pages. In essence, these navigation buttons control what is displayed in the main window. Each time you open the Double-Take Backup console, it will automatically open to the **Monitor** page. Any existing jobs will be displayed. If there are no existing jobs, the display will be blank.



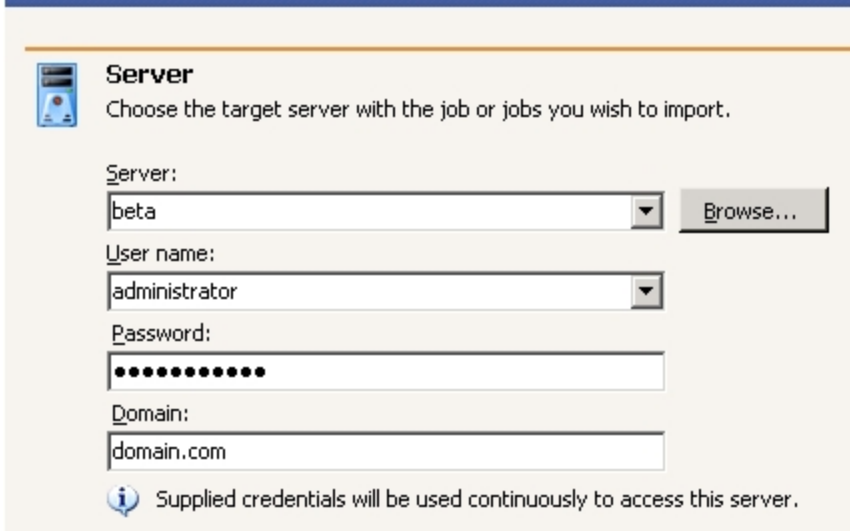
Note: The Double-Take Backup console requires a 1024x768 minimum screen resolution.

Importing a job

If you need to manage or monitor a job created on another console or created by another user, you can import it into your console session using the following instructions.

1. From the **Monitor** page, select **Import jobs** on the toolbar.
2. Identify the import server. This is the repository server that contains the source images that you are protecting.

Choose import server




Server
Choose the target server with the job or jobs you wish to import.

Server:
beta

User name:
administrator

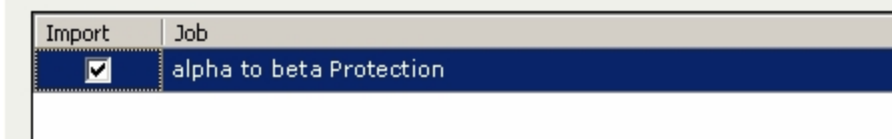
Password:
●●●●●●●●●●

Domain:
domain.com

 Supplied credentials will be used continuously to access this server.

- **Server**—To import a protection job (source to repository server), specify the name of the repository server. To import a recovery job (repository server to recovery server), specify the name of the recovery server. You can also click **Browse** to select the name from Active Directory or from a network drill-down list.
 - **User name**—Specify a user that is a member of the **Double-Take Admin** security group on the repository server.
 - **Password**—Specify the password associated with the **User name** you entered.
 - **Domain**—If you are working in a domain environment, specify the **Domain**.
3. Click **Next** to continue.
 4. Mark those jobs that you want to import into the console with a checkmark.

Choose jobs to import



Import	Job
<input checked="" type="checkbox"/>	alpha to beta Protection

5. Click **Finish** to complete the import.

The [Monitor](#) page will automatically appear with the imported job.

Clearing console security credentials

By default, the Double-Take Backup console caches security credentials. Also by default, the cache is saved when you close the console, allowing you to reuse the credentials when you reuse the console. This setting is controlled by the **Tools, Save credentials cache** menu option. When the option is selected with a checkmark, the security credentials cache will be saved. When the option is disabled and no checkmark is displayed, the security credentials cache will be deleted when the console is closed. At any time, you can clear the security credentials cache by selecting **Tools, Clear credentials cache**.

Source protection

When protecting a source, you can protect the entire server, including the system state, which is the server's operating system and configured applications, or you can protect only specific data on the server. Use the instructions [Protecting your source](#) for both full-server and data protection. At the end of the protection workflow, you will have an opportunity to configure [optional protection settings](#) for your protection job.

Protecting your source

Use the following instructions to set up your source protection.

1. Click **Protect** from the left navigation pane.
2. Identify the source server that you want to protect.

- **Server**—Specify the name of the source server. You can also click **Browse** to select the name from Active Directory or from a network drill-down list.
 - **User name**—Specify a user that is a member of the **Double-Take Admin** security group on the source.
 - **Password**—Specify the password associated with the **User name** you entered.
 - **Domain**—If you are working in a domain environment, specify the **Domain**.
3. Click **Next** to continue.
 4. Identify the repository server that will store the source data or the image of the entire source.

Choose repository server



Choose repository server

The repository server stores one or more server images for later recovery tasks.

Server:

beta

User name:


administrator

Password:

••••••••••

Domain:

domain.com

 Supplied credentials will be used continuously to access this server.

- **Server**—Specify the name of the repository server. You can also click **Browse** to select the name from Active Directory or from a network drill-down list.
 - **User name**—Specify a user that is a member of the **Double-Take Admin** security group on the repository server.
 - **Password**—Specify the password associated with the **User name** you entered.
 - **Domain**—If you are working in a domain environment, specify the **Domain**.
5. Click **Next** to continue.
 6. Select what you want to protect on the source.

Choose data to protect



Source data

Choose the source data to protect. System state file and folder cannot be modified.

- Protect system state.
- Protect selected data only. This data will not qualify for a full server recovery.

<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>		C:
<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>		E:
<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>		F:

Select all

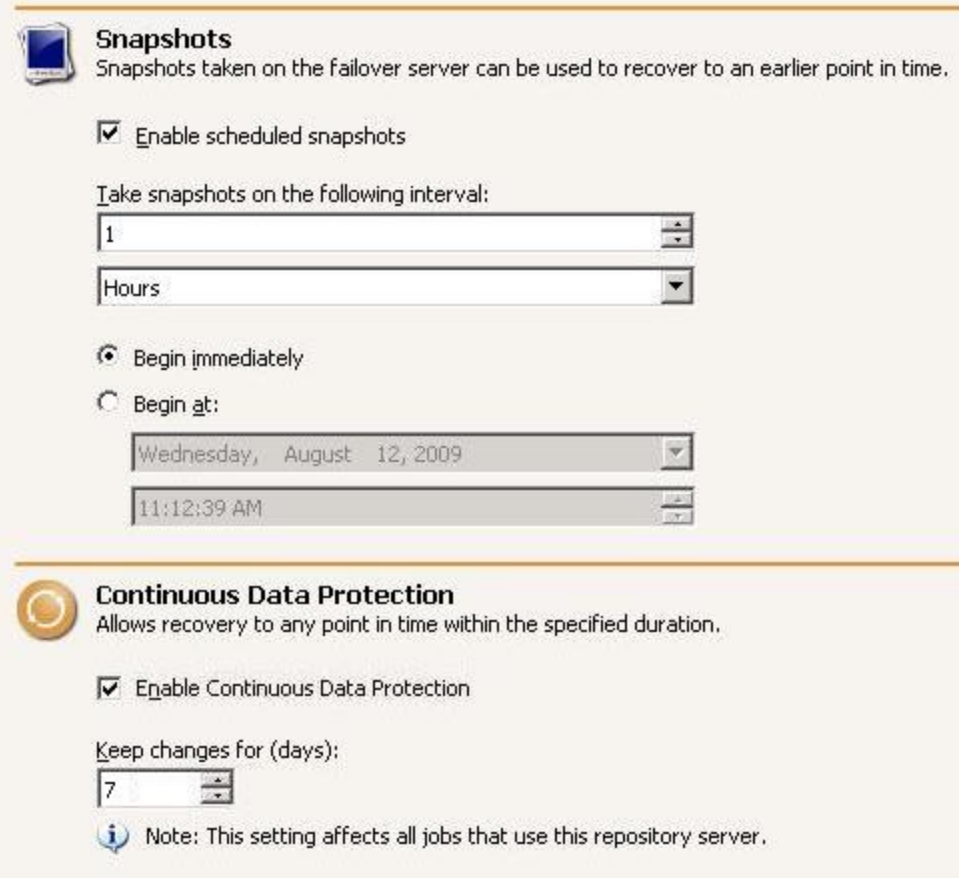
Deselect all

- **Protect system state**—This option protects the entire source, including the system state, which is the server's operating system and configured applications. If you select this option, you will be able to recover the entire source server.
 - **Protect selected data only**—This option protects only the data you select. If you select this option, you will only be able to recover the data you select but not the entire source server.
7. By default, Double-Take Backup selects all applicable source files based on your system state or selected data only choice. By default for system state protection, some files are included and excluded depending on if they can be used during a recovery. For example, the boot volume cannot be excluded because that is where the system state information is stored. If desired, you can exclude other files that you do not want to protect, but be careful when excluding data. Excluded volumes, folders, and/or files may compromise the integrity of your applications.
- Volumes, folders, and files that are marked with a checkmark are included. If there is no checkmark, the item is excluded. Expand and collapse the directory tree and click on an item to add or remove a checkmark.

Note: The **Select all** and **Deselect all** buttons allow you to quickly select or deselect all data on your server based on your system state or selected data only choice. These buttons will do not change what is selected for the system state, which is defined by Double-Take Backup and cannot be modified.

- Once you have configured the data to protect, click **Next** to continue.
- A point-in-time strategy gives you additional recovery options. In addition to being able to recover from live data, you can recover to a single point in time using a snapshot or to any point in time using TimeData continuous data protection.

Choose point-in-time strategy



Snapshots
Snapshots taken on the failover server can be used to recover to an earlier point in time.

Enable scheduled snapshots

Take snapshots on the following interval:

1

Hours

Begin immediately

Begin at:

Wednesday, August 12, 2009

11:12:39 AM

Continuous Data Protection
Allows recovery to any point in time within the specified duration.

Enable Continuous Data Protection

Keep changes for (days):

7

Note: This setting affects all jobs that use this repository server.

- **Enable scheduled snapshots**—By default, periodic snapshots are disabled. If desired, enable snapshots by selecting this option. Because Double-Take Backup uses the Microsoft Volume Shadow Copy service to create snapshots, your source and repository server must be using the NTFS file system. If you are using a FAT file system, the FAT volumes will not be included in the snapshot set, and when the snapshots are reverted, the FAT volume will not be time-consistent with the NTFS volumes.
- **Take snapshots on the following interval**—By default, Double-Take Backup will take a snapshot of the repository server data every hour. If desired, increase or decrease the interval between snapshots.
- **Begin immediately**—If you want to start taking snapshots immediately after the Double-Take Backup job is established, select this option.
- **Begin at**—If you want to start taking snapshots at a specific date and time, select this option and specify the date and time parameters.
- **Enable Continuous Data Protection**—By default, continuous data protection is disabled. If desired, enable continuous data protection by selecting this option. If you did

not install TimeData, you will not have this option available to select. The repository server must be using the NTFS file system. If you are using a FAT file system, the FAT volumes will not be included in the continuous data protection and will be unavailable for point-in-time recovery.

- **Keep changes for**—By default, Double-Take Backup will capture and catalog all changes made to all files and folders for seven days. The recommended retention is seven days. If desired, increase or decrease the number of days to protect data continuously. If you increase the number of days, keep in mind [space limitations on the repository server](#). Continuous data protection is a per repository server setting. Therefore, the length of time you specify will apply to all jobs that use this repository server.
9. Click **Next** to continue.
 10. Select a location on the repository server to store the location of the source's data or image. The location you specify for **Path prefix** will automatically have the source server name appended to it, in order to distinguish the data or image from other data or images.

Choose image storage location

Image storage location
Select a location on the repository server to store this image. The source name will automatically be appended to the path.

Path prefix:
C:\Server Image

Full Path:
C:\Server Image\alpha

Delete Orphans

11. Specify if you want to delete orphan files. An orphan file is a file that exists in the repository server's copy of the data or source image, but it does not exist on the source. An orphan file can be created when a job is stopped and file is deleted on the source. Because of the stopped job, the delete operation does not get replicated to the repository server. Orphan files may also exist if files were manually copied into the repository server's copy of the data or source image.
12. Click **Next** to continue.
13. At the **Protection Summary** page, you can review the choices you have made and you can also set optional data transmission and e-mail notification settings. If you need to make any changes to the selections you made, click **Back** to return to the previous pages. If you want to [configure the optional settings](#), click **Configure**.
14. Once you have completed your configuration, click **Finish** to establish your source protection. The [Monitor](#) page will automatically appear with the new job.

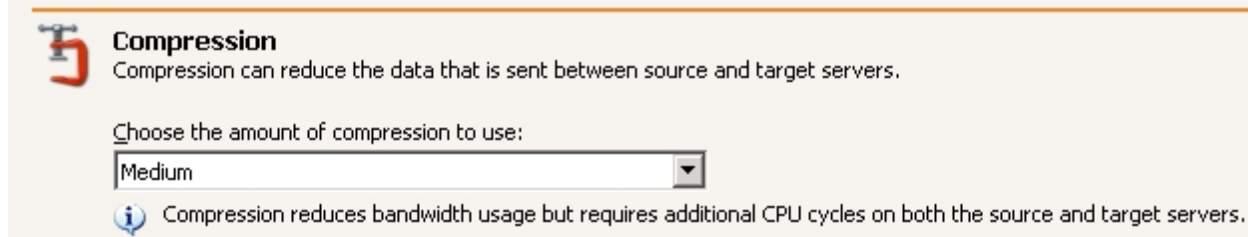
Optional protection settings

When you are creating a protection job, you will have an opportunity to configure optional job settings when you get to the **Protection summary** workflow page. If you want to configure optional settings, click **Configure** next to the options section.

- [Compressing data for transmission for a protection job](#)
- [Limiting bandwidth for a protection job](#)
- [Routing transmission for a protection job](#)
- [E-mailing repository server event messages](#)


Compressing data for transmission for a protection job

To help reduce the amount of bandwidth needed to transmit Double-Take Backup 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 repository server receives the compressed data, it decompresses it and then writes it to disk.



Compression
Compression can reduce the data that is sent between source and target servers.

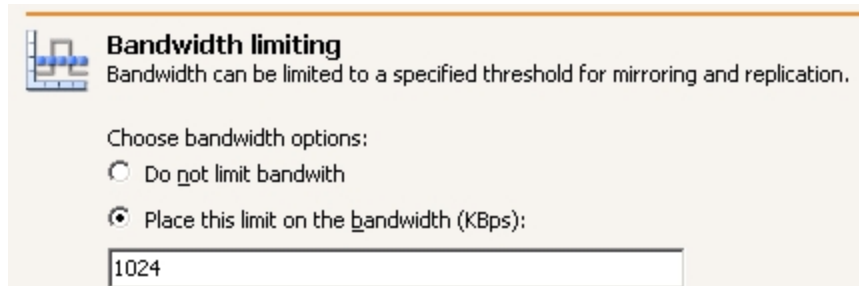
Choose the amount of compression to use:
Medium

 Compression reduces bandwidth usage but requires additional CPU cycles on both the source and target servers.

By default, compression is set to **None**, which means it is disabled. The compression levels can be set at **Minimum**, **Medium**, or **Maximum**. Each level requires more overhead than the previous level of compression. 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.

Limiting bandwidth for a protection job

Bandwidth limitations are available to restrict the amount of network bandwidth used for Double-Take Backup data transmissions. When a bandwidth limit is specified, Double-Take Backup never exceeds that allotted amount. The bandwidth not in use by Double-Take Backup is available for all other network traffic.



Bandwidth limiting
Bandwidth can be limited to a specified threshold for mirroring and replication.

Choose bandwidth options:

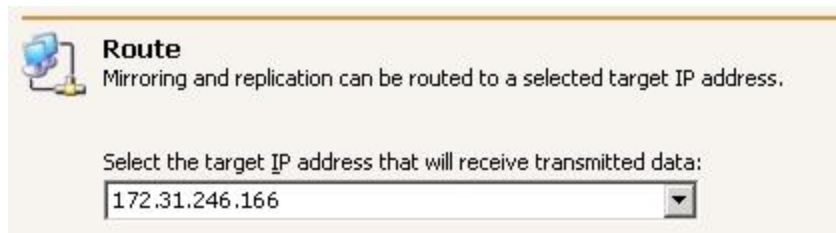
- Do not limit bandwidth
- Place this limit on the bandwidth (KBps):

1024

- **Do not limit bandwidth**—Double-Take Backup will transmit data using 100% bandwidth availability.
- **Place this limit on the bandwidth**—Enter a value in kilobytes per second to limit data transmission. This is the maximum amount of data that will be transmitted per second.

Routing transmission for a protection job

By default, Double-Take Backup will select a default route for transmissions. The default route is the first connection listed in the NIC binding order on the repository server. If you want to select a different route for Double-Take Backup transmissions, you can select a different IP address on the repository server using the Double-Take Backup **Route** option, or you can reorder the bindings on the repository server (**Network Connections, Advanced, Advanced Settings, Adapters and Bindings**).



Route
Mirroring and replication can be routed to a selected target IP address.

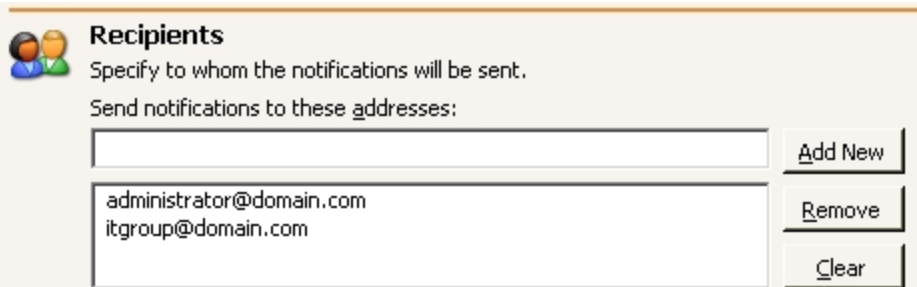
Select the target IP address that will receive transmitted data:
172.31.246.166

E-mailing repository server event messages

E-mail notification sends specific Double-Take Backup events to user-defined e-mail addresses. If you want to use e-mail notification, you will need to identify who will receive the messages, what messages they will receive, and your e-mail server configuration.

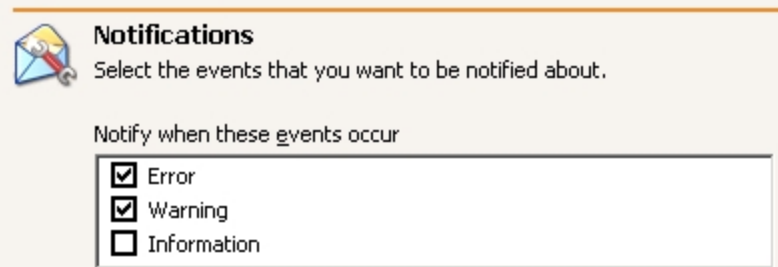
Note: The e-mail notification settings are for the repository server. If you want to configure e-mail notification for your source server, see [E-mailing source event messages](#).

- **Recipients**—Specify the e-mail address that each Double-Take Backup e-mail message should be sent to and click **Add New**. The e-mail address will be inserted into the list of addresses. Each address is limited to 256 characters, and 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**. If you want to remove all of the addresses from the list, click **Clear**.



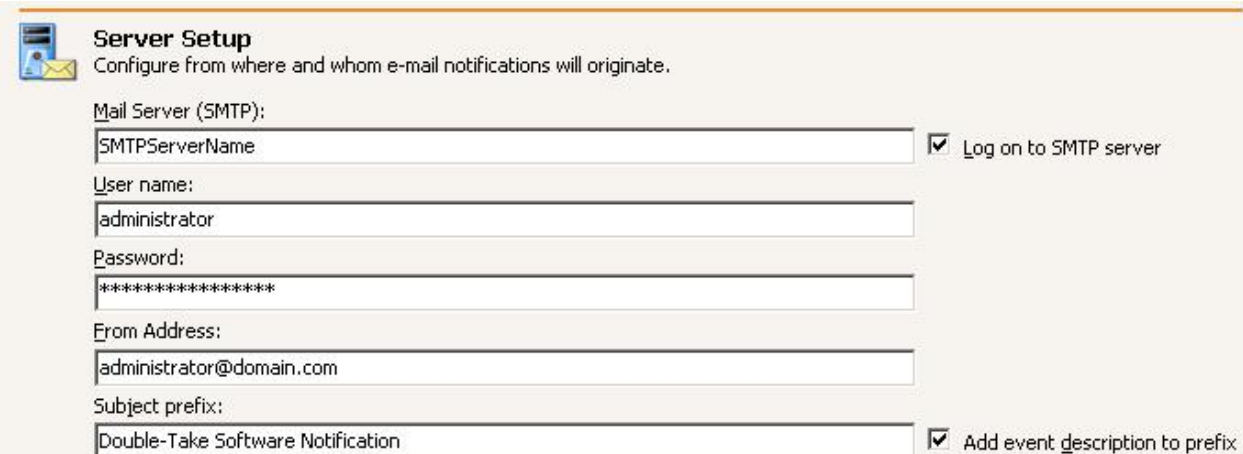
The screenshot shows the 'Recipients' configuration panel. It has a title 'Recipients' with a group of three people icon. Below the title is the instruction 'Specify to whom the notifications will be sent.' and 'Send notifications to these addresses:'. There is a text input field for adding new addresses, followed by a list of existing addresses: 'administrator@domain.com' and 'itgroup@domain.com'. To the right of the list are three buttons: 'Add New', 'Remove', and 'Clear'.

- **Notifications**—Specify which events that you want to be sent via e-mail. You can select **Error**, **Warning**, and/or **Information**.



The screenshot shows the 'Notifications' configuration panel. It has a title 'Notifications' with a mail icon. Below the title is the instruction 'Select the events that you want to be notified about.' and 'Notify when these events occur'. There is a list of three events with checkboxes: 'Error' (checked), 'Warning' (checked), and 'Information' (unchecked).

- **Server Setup**—Specify your e-mail server configuration.



The screenshot shows the 'Server Setup' configuration panel. It has a title 'Server Setup' with a mail icon. Below the title is the instruction 'Configure from where and whom e-mail notifications will originate.' There are several fields: 'Mail Server (SMTP):' with a text input field containing 'SMTPServerName' and a checked checkbox 'Log on to SMTP server'; 'User name:' with a text input field containing 'administrator'; 'Password:' with a text input field containing '*****'; 'From Address:' with a text input field containing 'administrator@domain.com'; and 'Subject prefix:' with a text input field containing 'Double-Take Software Notification' and a checked checkbox 'Add event description to prefix'.

- **Mail Server**—Specify the name of your SMTP mail server.
- **Log on to SMTP server**—If your SMTP server requires a login, enable this option and specify a **User name** and **Password**.

Note: TimeData events do not support authenticated SMTP notification methods. Check the [TimeData log files](#) for TimeData events if you are using authenticated SMTP.

- **User name**—Specify a user account with privileges to send e-mail messages from your SMTP server.
 - **Password**—Specify the password associated with the **User Name** you entered.
 - **From Address**—Specify the e-mail address that you want to appear in the **From** field of each Double-Take Backup e-mail message. The address is limited to 256 characters.
 - **Subject prefix**—Specify default text that will appear on every Double-Take Backup e-mail message. This will help distinguish Double-Take Backup messages from other messages.
 - **Add event description to prefix**—Specify if you want the description of the event appended to the **Subject prefix**. The subject line is limited to 150 characters, so only the first 150 characters of the combined **Subject prefix** and event description will be displayed.
-

Note: Make sure you configure your e-mail client so that the Double-Take Backup e-mail notification messages do not get blocked by a spam filter or sent to a junk mail folder.

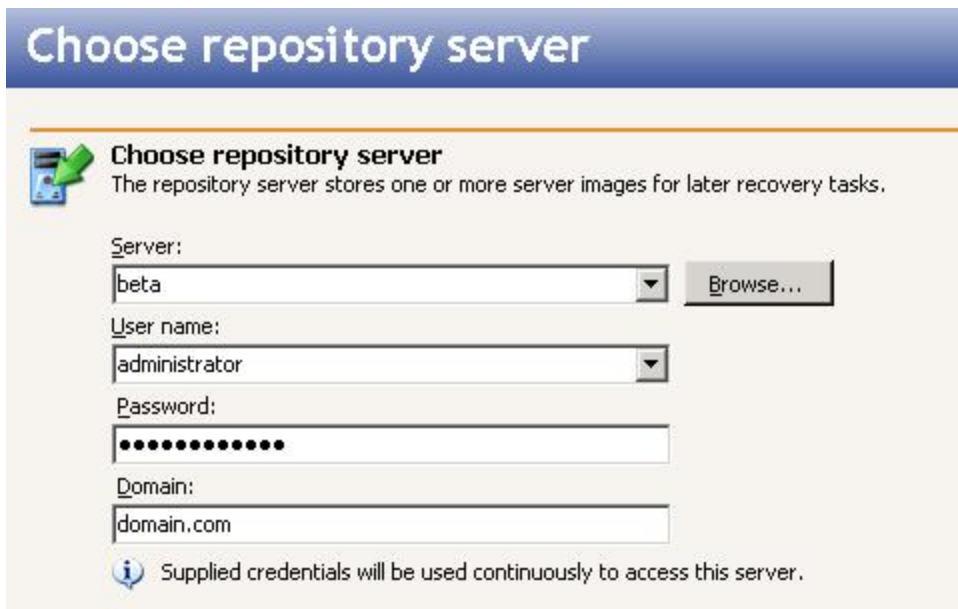
Data on the repository server

The images that are stored on the repository server may be in various states. There is live data, snapshot data, and TimeData historical data. There may be times when you want to view or retrieve data from the repository server without recovering an entire server or all of the data you have protected.

- **Live data**—The live data can be easily be viewed and retrieved at any time on the repository server using standard file system tools like Windows Explorer. Be careful not to modify any of the data on the repository server, so that if the data is needed for a recovery, it will be the same as what is on the source.
- **Snapshot and TimeData historical data**—Because the snapshot data and TimeData historical data is not in client accessible format, you need to use Double-Take Backup to [mount a snapshot or a TimeData point in time](#) so that it can be accessed. Once mounted, you can retrieve files or folders as needed using standard file system tools like Windows Explorer.
- **Exchange data**—If you have installed Ontrack PowerControls, you can open a Microsoft Exchange storage file (.edb file) for [viewing or retrieving individual messages or mailboxes](#). Ontrack PowerControls can be used on the live data, a snapshot, or TimeData historical data.

Viewing and retrieving data from snapshots or from a TimeData point in time

1. Click **View Point In Time** from the left navigation pane.
2. Identify the repository server that is storing the source image(s).



- **Server**—Specify the name of the repository server. You can also click **Browse** to select the name from Active Directory or from a network drill-down list.
 - **User name**—Specify a user that is a member of the **Double-Take Admin** security group on the repository server.
 - **Password**—Specify the password associated with the **User name** you entered.
 - **Domain**—If you are working in a domain environment, specify the **Domain**.
3. Click **Next** to continue.
 4. Select the image that you want to view and specify the point-in-time strategy.

Choose image



Source server image

Select an image containing point-in-time data

Available images:

Image Name	Image Location	Image State	Browse...
alpha	c:/server image/alpha	Good	



Point-in-time strategy

Snapshot

Date	Type	Status
7/23/2009 12:53:55 PM	User Request	Ok

Specific time

Data is valid for recovery during these times:

Start Time	End Time
7/23/2009 12:53:55 AM	7/25/2009 1:04:39 PM
7/25/2009 1:07:02 PM	7/31/2009 2:17:21 PM
7/31/2009 3:58:24 PM	8/12/2009 11:44:49 AM

- **Available images**—Select the image of the source that contains the data you want to recover. If the image is not displayed, click **Browse** to look for the location of the server image.
 - **Snapshot**—Select this option and choose a snapshot from the list to mount the data stored in that snapshot. If you did not enable snapshots or there were no snapshots taken on the repository server, there will be no snapshots in the list.
 - **Specific time**—Select this option and choose a specific date and time to mount the data from that point in time. The table below the date and time fields allows you to confirm that the Double-Take Backup data is valid for the point in time you are selecting. If you did not enable continuous data protection when you configured protection or if the TimeData initialization has not completed, this option will not be available.
5. Click **Finish**. If you are completing this task from the repository server, the data from the image and point-in-time strategy you selected will be mounted automatically and a Windows Explorer window will open to the mount point. If you are not completing this task from the repository server, the data from the image and point-in-time strategy you selected will still be mounted automatically, but you will see a pop-up dialog box indicating where you can find the mount point on the repository server.
 6. From Windows Explorer, retrieve files or folders as needed. If the files or folders you need are not in the snapshot or TimeData point in time, you can repeat the process multiple times to mount different snapshots or TimeData points in time.

Viewing and retrieving Exchange data

1. Start Ontrack PowerControls by selecting **Start, Programs, Kroll Ontrack, Ontrack PowerControls for Double-Take, Ontrack PowerControls for Exchange**.
2. The Data Wizard will automatically open the first time Ontrack PowerControls is opened. If it is not opened, select **File, Use Wizard**.
3. Specify the storage files from your repository server that you want to browse. If you want to browse a snapshot or TimeData historical data, make sure you [mount the snapshot or TimeData historical data](#) first.
4. Click **Next** to continue.
5. Specify where you want to copy the Exchange data.
6. Click **Next** to continue.
7. After Ontrack PowerControls has scanned the files, click **Finish**.

For more detailed information on using Ontrack PowerControls and its other features, press F1 from within Ontrack PowerControls to open the online help.

Recovery

In some cases, you can recover data or an entire server back to your original source. Additionally, you can use the source data or the image of a entire source, stored on the repository server, to quickly and easily create a new source server. Your new source can be an existing physical or virtual machine, or Double-Take Backup can automatically provision (create) a virtual machine during the recovery process.

When your recovery server is a not your original source, Double-Take Backup must be installed on the recovery server (or on the existing virtual machine on an ESX server) before you start the recovery process. However, you do not need to install Double-Take Backup from a CD or web download or have a valid activation code for the recovery server. Double-Take Backup has a built-in installation feature that installs a valid, unactivated copy of Double-Take Backup on the recovery server. At any time before the recovery, you can use this built-in installation feature to push a copy of Double-Take Backup to your recovery server. If your recovery server already has a previously installed, properly licensed and activated copy of Double-Take Backup, you can skip the recovery installation and go directly to the recovery process. See [Installing on the recovery server](#) for the steps to install Double-Take Backup on your recovery server.

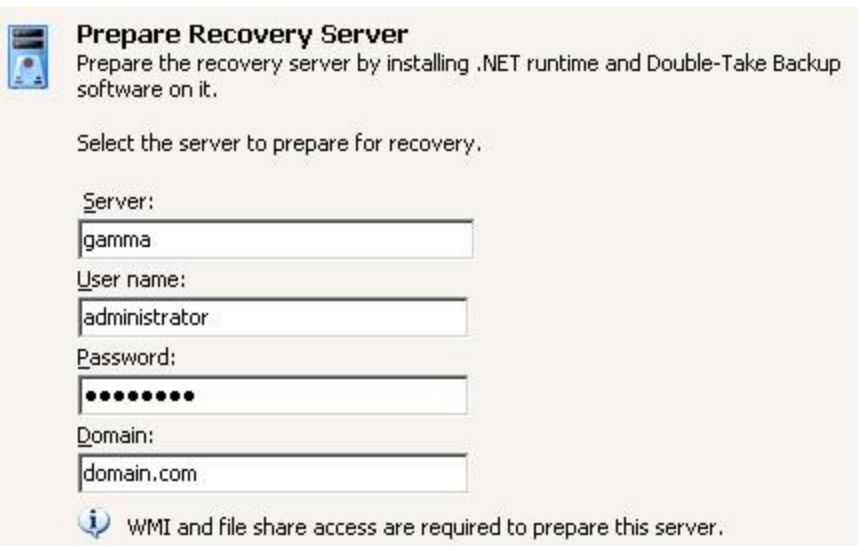
Before you begin the recovery process, understand that the flow of data is going to change. When you were protecting your source, the data was being transmitted from the source to the repository server. During recovery, the data will be transmitted from the repository server to the recovery server. This means your repository server is now your source, and your recovery server is your target.

You have the following recovery options.

- [Recovering an entire server to an existing physical or virtual machine](#)
- [Recovering an entire server to an automatically provisioned virtual machine on VMware ESX](#)
- [Recovering an entire server to an automatically provisioned virtual machine on Hyper-V](#)
- [Recovering data to an existing physical or virtual machine](#)
- [Optional recovery settings](#)

Installing Double-Take Backup on the recovery server

1. Select **Tools, Prepare Recovery Server**.
2. Specify the recovery server information.



Prepare Recovery Server
Prepare the recovery server by installing .NET runtime and Double-Take Backup software on it.

Select the server to prepare for recovery.

Server:
gamma

User name:
administrator

Password:
••••••••

Domain:
domain.com

WMI and file share access are required to prepare this server.


- **Server**—Specify the name of the recovery server.
 - **User name**—Specify a user that has WMI and file share access on the server.
 - **Password**—Specify the password associated with the **User name** you entered.
 - **Domain**—If you are working in a domain environment, specify the **Domain**.
3. Click **Install** and monitor the status of the various installation tasks.

Note: Do not use the copy of Double-Take Backup that was pushed to your recovery server outside of the recovery process because it will automatically activate the software. Once activated, it will automatically deactivate after five days.

Recovering an entire server to an existing physical or virtual machine


1. There are two ways to begin the recovery process.
 - Click **Recover** from the left navigation pane and identify the repository server that contains your source data that you want to recover. Click **Next** to continue.
 - Highlight the data protection job on the **Monitor** page and select **Recover** from the toolbar. This automatically identifies the repository server that contains the source data that you want to recover.
2. Select the image of the source that you want to recover and the point-in-time strategy.

Choose image

 **Source server image**
Select an image

Available images:

Image Name	Image Location	Image State	Browse...
alpha	c:/server image/alpha	Good	

 **Point-in-time strategy**

Live data

Snapshot

Date	Type	Status
7/23/2009 12:53:55 PM	User Request	Ok

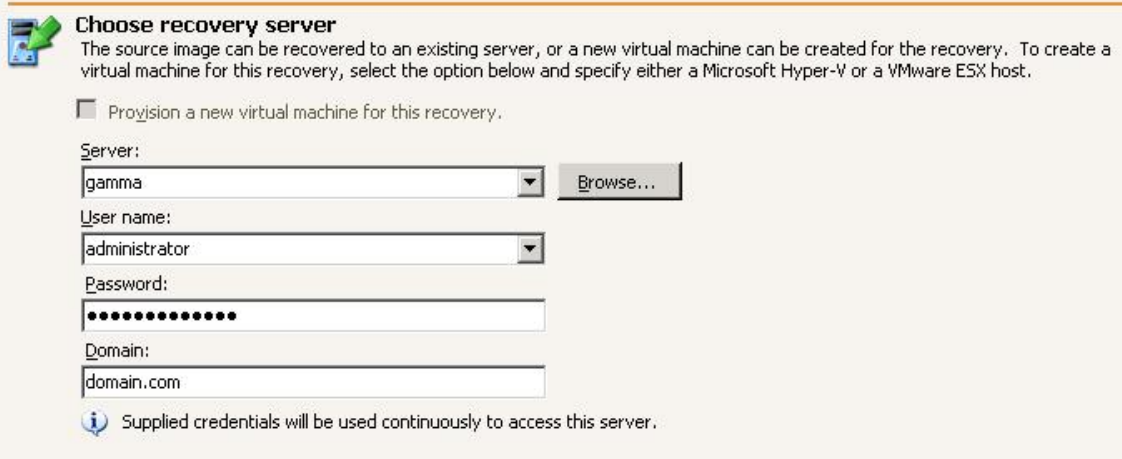
Specific time

Data is valid for recovery during these times:

Start Time	End Time
7/23/2009 12:53:55 AM	7/25/2009 1:04:39 PM
7/25/2009 1:07:02 PM	7/31/2009 2:17:21 PM
7/31/2009 3:58:24 PM	8/12/2009 11:44:49 AM

- **Available images**—Select the image of the source that contains the data you want to recover. If the image is not displayed, click **Browse** to look for the location of the server image.
 - **Live data**—Select this option to recover to the live data that is currently stored on the repository server.
 - **Snapshot**—Select this option and choose a snapshot from the list to recover to the data stored in that snapshot. If you did not enable snapshots or there were no snapshots taken on the repository server, there will be no snapshots in the list.
 - **Specific time**—Select this option and choose a specific date and time to recover the data from that point in time. The table below the date and time fields allows you to confirm that the Double-Take Backup data is valid for the point in time you are selecting. If you did not enable continuous data protection when you configured protection or if the TimeData initialization has not completed, this option will not be available.
3. Click **Next** to continue.
 4. Select the recovery server. If you used the Double-Take Backup built-in installation to push a copy of Double-Take Backup to the recovery server, that copy will automatically be activated when you select it as a recovery server. The copy will automatically deactivate after five days.

Choose recovery server



Choose recovery server

The source image can be recovered to an existing server, or a new virtual machine can be created for the recovery. To create a virtual machine for this recovery, select the option below and specify either a Microsoft Hyper-V or a VMware ESX host.


Provision a new virtual machine for this recovery.

Server:

User name:

Password:

Domain:

 Supplied credentials will be used continuously to access this server.

- **Provision a new virtual machine for this recovery**—Do not enable this option. This setting is only if you are recovering to an automatically provisioned virtual machine.
 - **Server**—Specify the name of the recovery server. You can also click **Browse** to select the name from Active Directory or from a network drill-down list. You can select your repository server as the recovery server, but keep in mind that the repository server role and any other source images will be lost.
 - **User name**—Specify a user that is a member of the **Double-Take Admin** security group on the recovery server. If your original source was the only domain controller in your network, you must specify a local account in the local administrators group.
 - **Password**—Specify the password associated with the **User name** you entered.
 - **Domain**—If you are working in a domain environment, specify the **Domain**.
5. Click **Next** to continue.
 6. By default, Double-Take Backup will **Recover the entire server**. If you only want to recover selected data, select **Recover selected data only**. In either case, you can, if desired, exclude files that you do not want to recover, but be careful when excluding data. Excluded volumes, folders, and/or files may compromise the integrity of your applications.

Choose image data to recover



Source data

Choose the source data to recover. System state file and folder cannot be modified.

- Recover the entire server.
- Recover selected data only. This will not perform a full server recovery.



Volumes, folders, and files that are marked with a checkmark are included. If there is no checkmark, the item is excluded. Expand and collapse the directory tree and click on an item to add or remove a checkmark. Once you have configured the data to recover, click **Next** to continue.

Note: The **Select all** and **Deselect all** buttons allow you to quickly select or deselect all data on your server based on your system state or selected data only choice. These buttons will do not change what is selected for the system state, which is defined by Double-Take Backup and cannot be modified.

7. Specify if your recovery server is in a LAN or WAN environment with respect to your original source.

Set recovery type



Recovery type

Identify if your original source and your recovery server are in a LAN or WAN environment.

LAN recovery

For a LAN recovery, the network identity of the original source will be applied to the recovery server. The mapping between original source NICs and recovery server NICs must be configured.

Map network adapters:

Source Network Adapter	Target Network Adapter
Public network	Public network
Private network	Public network

WAN recovery

For a WAN recovery, the name of the source server is applied to the recovery server, however the IP addresses of the recovery server are used and DNS updates are made to point the server name to the new IP addresses. Additional IP addresses can be added.

Update DNS Servers

- **LAN recovery**—If your recovery server is on a LAN with your original source, the original source name and the IP address(es) will be applied to the recovery server. Specify which network adapters on the recovery server that you want to take over for the network adapters on the original source.
 - **WAN recovery**—If your recovery server is located across a WAN from your original source, the original source name will be applied to the recovery server, but the IP address(es) of the original source will not be used. Instead, the IP address(es) of the recovery server will be used. If desired, you can select to **Update DNS servers**, which will allow you to specify how to resolve server names and IP addresses. If you do not select **Update DNS Servers**, you must manually modify DNS after the recovery is complete.
8. Click **Next** to continue.
 9. If you selected to update your DNS servers for a WAN recovery type, specify your DNS update options.
 - **DNS Domains**—Click **Add** and specify the name of your DNS Domain. Repeat this for each DNS domain. If you want to remove a domain in the list, highlight the name and click **Remove**.
 - **User name**—Highlight a DNS domain and specify a user account that has privileges to update DNS.
 - **Password**—Specify the password associated with the user account you entered.
 - **DNS Servers**—Click **Add** and specify the IP address of a DNS server in the DNS domain. Repeat this for each DNS server in the DNS domain. If you want to remove a server from the list, highlight the address and click **Remove**.
 - **IP addresses**—For each IP address on the source, select an IP address on the recovery server that will take over for that source IP address.
 10. Once you have your DNS updates configured, click **Test DNS** to determine if your DNS updates will be successful.
 11. Click **Next** to continue.
 12. Specify your recovery shutdown options.

Set recovery shutdown options



Manual or Automatic Recovery

The recovery process can continue automatically after a mirror completes, or it can replicate changes while waiting for user intervention.

- Wait for user intervention before recovery



Shutdown Source Server

If the source server is detected on the network, it can be shutdown before the final reboot of the server to avoid identity conflicts.

- Shut down the source server



Original Protection Job

Any updates to the image during recovery could corrupt the recovered server.

- If the original source server is already unavailable, this choice will not be observed.
- Delete the original protection job.
- Stop the original protection job. The job can be restarted after recovery.
- Leave the original protection job running. The recovered server could be corrupted, however the original source will remain protected.

- **Wait for user intervention before recovery**—Specify if you want to pause the recovery process after the original source image has been mirrored to the recovery server, which allows you time to coordinate any other tasks. If you do not pause the recovery, the original source system state will be processed on the recovery server immediately after the mirror is complete. (The recovery server will automatically reboot after the system state processing is completed.)

-
- **Shut down the source server**—Specify if you want to shutdown the source server before the original source system state is processed on the recovery server, which avoids identity conflicts on the network. The shutdown will occur after the original source image is mirrored to the recovery server but before the original source system state is applied to the recovery server. This option is not available if the recovery server is the original source.
 - **Original Protection Job**—Specify what to do with the original protection job if the original source is still online. If the original source is not running, these options will be discarded because they cannot be performed. These options are not available if the recovery server is the original source or the repository server.
 - **Delete the original protection job**—The original protection job is stopped and then deleted.
 - **Stop the original protection job**—The original protection job is stopped but is not deleted. It can be restarted after the recovery.
 - **Leave the original protection job running**—The original protection job is not stopped or deleted. Keep in mind with this option that the data on the recovery server could become inconsistent or corrupted because the original source protection job could be transmitting data to the repository server and that data might get transmitted to the recovery server, depending on the progress of the recovery process .
13. Click **Next** to continue.
 14. At this point, Double-Take Backup validates that your recovery server is compatible to become your original source. Errors are designated by a red circle with a white X. (A question mark icon is an unknown error.) Warnings are designated by a yellow triangle with a black exclamation point. A successful validation is designated by a green circle with a white checkmark. You must correct any errors before you can start the recovery. You must revalidate the selected server until the validation check passes without errors. Once the validation is successful, click **Next** to continue.
 15. At the **Recovery summary** page, you can review the choices you have made and you can also set optional data transmission and e-mail notification settings. If you need to make any changes to the selections you made, click **Back** to return to the previous pages. If you want to [configure optional settings](#), click **Configure**. Once you have completed your configuration, click **Finish** to start the recovery.
 16. Monitor the progress of the recovery mirror on the [Monitor](#) page.
 17. When the recovery mirror has completed, the recovery server will automatically reboot to complete the recovery process. If you specified to pause the recovery before applying the system state, when the mirror is complete, the **Activity** will change to **Protected**. Use this time to complete any necessary tasks. When you are ready to complete the recovery, click **Recover** on the **Monitor** page toolbar to complete the recovery process. Once the recovery server has been rebooted, it will become the original source server.

Note Because the Windows product activation is dependent on hardware, you may need to reactivate your Windows registration after recovery. Follow the on-screen prompts to complete the reactivation.

Recovering an entire server to an automatically provisioned virtual machine on VMware ESX

If you want Double-Take Backup to create a virtual machine on VMware ESX during the recovery process, you must have a VMware ESX host machine with an existing virtual machine. See [Recovery server requirements](#) for details on the requirements for the ESX machine and the existing virtual machine. If your environment does not meet those requirements, you will have to recover to an existing physical or virtual machine.

The existing virtual machine used during the provisioning process is used to as an intermediary during the recovery process to create the new virtual server that, once online, will have the identity,

data, and system state of the original source. The existing virtual must have both Windows and [Double-Take Backup installed and licensed](#) before you can start the recovery.

Tasks performed by the existing virtual machine during the recovery process

1. Create a new virtual machine
2. Add the disk(s) for that virtual machine to its own machine
3. Mount the disk(s)
4. Apply the incoming mirror data (the original source data and system state information) from the repository server to the mounted disk(s)
5. Unmount the disk(s)
6. Remove the disk(s) from its own machine
7. Start the new virtual machine

Once the new virtual machine is online, it will have the identity, data, and system state of the original source. Since the existing virtual machine maintains its own identity, it can be reused for additional recoveries.

Recovery steps

1. There are two ways to begin the recovery process.
 - Click **Recover** from the left navigation pane and identify the repository server that contains your source data that you want to recover. Click **Next** to continue.
 - Highlight the data protection job on the **Monitor** page and select **Recover** from the toolbar. This automatically identifies the repository server that contains the source data that you want to recover.
2. Select the image of the source that you want to recover and the point-in-time strategy.

Choose image



Source server image

Select an image

Available images:

Image Name	Image Location	Image State	Browse...
alpha	c:/server image/alpha	Good	



Point-in-time strategy

Live data

Snapshot

Date	Type	Status
7/23/2009 12:53:55 PM	User Request	Ok

Specific time

Data is valid for recovery during these times:

Start Time	End Time
7/23/2009 12:53:55 AM	7/25/2009 1:04:39 PM
7/25/2009 1:07:02 PM	7/31/2009 2:17:21 PM
7/31/2009 3:58:24 PM	8/12/2009 11:44:49 AM

- **Available images**—Select the image of the source that contains the data you want to recover. If the image is not displayed, click **Browse** to look for the location of the server image.
 - **Live data**—Select this option to recover to the live data that is currently stored on the repository server.
 - **Snapshot**—Select this option and choose a snapshot from the list to recover to the data stored in that snapshot. If you did not enable snapshots or there were no snapshots taken on the repository server, there will be no snapshots in the list.
 - **Specific time**—Select this option and choose a specific date and time to recover the data from that point in time. The table below the date and time fields allows you to confirm that the Double-Take Backup data is valid for the point in time you are selecting. If you did not enable continuous data protection when you configured protection or if the TimeData initialization has not completed, this option will not be available.
3. Click **Next** to continue.
 4. Select the recovery server. This is the ESX server that will host the virtual machine that, after the recovery, will become your new source.

Choose recovery server



Choose recovery server

The source image can be recovered to an existing server, or a new virtual machine can be created for the recovery. To create a virtual machine for this recovery, select the option below and specify either a Microsoft Hyper-V or a VMware ESX host.

Provision a new virtual machine for this recovery.

Server:

User name:

Password:

Domain:

Supplied credentials will be used continuously to access this server.
The domain entry will be ignored for a VMware ESX host.

- **Provision a new virtual machine for this recovery**—Enable this option to recover to an automatically provisioned virtual machine hosted on an ESX server.
 - **Server**—Specify the name of the ESX server. You can also click **Browse** to select the name from Active Directory or from a network drill-down list.
 - **User name**—Specify the root user or another user that has the administrator role on the VMware ESX server.
 - **Password**—Specify the password associated with the **User name** you entered.
 - **Domain**—This option is not necessary for an ESX server.
5. Click **Next** to continue.
 6. Because you are creating a new virtual machine, you must **Recover the entire server** so that the system state from the original source can be used to create the new virtual. You can exclude volumes that you protected but do not want to recover, but be careful when excluding data. Excluded volumes may compromise the integrity of your applications. Volumes marked with a checkmark are included. If there is no checkmark, the item is excluded.

Choose image data to recover



Source data

Choose the source data to recover. System state file and folder cannot be modified.

- Recover the entire server.
- Recover selected data only. This will not perform a full server recovery.

- | | | |
|-------------------------------------|--|----|
| <input checked="" type="checkbox"/> | | C: |
| <input checked="" type="checkbox"/> | | E: |
| <input checked="" type="checkbox"/> | | F: |

7. Click **Next** to continue.
8. Specify the volumes to create on the new virtual machine.

Edit provisioned volumes



Provisioned Volumes

Edit the provisioned volumes on the virtual machine that will take over for the source server after recovery.

Volume	Selected Data (GB)	Source Size (GB)	Provisioned Volume (GB)
C	5.91	9.99	9.99
E	0.10	4.99	4.99
F	218.68	512.65	512.65

- **Selected Data**—This displays the amount of data, in gigabytes, currently on the source volume.
 - **Source Size**—This displays the size, in gigabytes, of the volume on the source.
 - **Provisioned Volume**—Specify the size, in gigabytes, of the volume to be created on the new virtual.
9. Click **Next** to continue.
 10. Choose a datastore on the recovery server to store the protected data and system state information that will be incoming from the repository server. You can only select a datastore that has enough free space available. The amount of necessary disk space is noted at the bottom of the page.

Choose Datastore



Target datastore

The protected data can be stored on any datastore that the target server can access.

Choose a target datastore:

Datastore	Total Size	Free Space
local 1	26.25 GB	9.63 GB
local 2	546.25 GB	202.13 GB

Note: If the size of the datastore is identical to the size of the disk on original source and there is less than 20 MB of free space on that original source, you may run out of disk space on the datastore due to differences in how the virtual disk block size is formatted. In this case, make sure that your datastore has at least 20 MB more disk space than the size of the disk on the original source.

11. Click **Next** to continue.
12. Choose the Virtual Recovery Appliance, which is your existing virtual machine, that will be used to create the new virtual machine. The existing virtual machine must have both Windows and Double-Take Backup installed and licensed on it.

Choose Virtual Recovery Appliance



Choose the Virtual Recovery Appliance

The Virtual Recovery Appliance hosts the datastore that will ultimately become the original source. Double-Take Backup must be installed and running on the machine.

Server:

gamma

User name:

administrator

Password:

•••

Domain:



Supplied credentials will be used continuously to access this server.



VMWare Tools

The latest VMWare Tools needs to be installed in order to use the following servers.

No servers were located with dated VMWare Tools.

- **Server**—Specify the name of the existing virtual machine on the ESX server. The drop-down list will only display virtual machines that are running and have an operating system capable of performing a recovery.
 - **User name**—Specify a user that is a member of the **Double-Take Admin** security group on the existing virtual machine. The user must also be a member of the local administrators group. If your original source was the only domain controller in your network, the user must be a local account in the local administrators group.
 - **Password**—Specify the password associated with the **User name** you entered.
 - **Domain**—If you are working in a domain environment, specify the **Domain**.
 - **VMware Tools**—Any servers from the **Server** list that do not have the latest VMware tools installed will be listed.
13. Click **Next** to continue.
 14. Configure the new virtual machine that will be created and will become the source after the recovery.

Edit replica settings



Replica virtual machine

Configure the virtual machine that will take over for the source server after cutover.

VMware Console display name

Map replica virtual network IP to target networks:

Replica VM Network IP	Target Network
172.31.246.166	VM Network

Number of processors:

Processors on the source server:

Amount of memory (MB):

Memory on the source server (MB):

- **VMware Console display name**—By default, source name_VM will be used to create the new virtual machine display name. If desired, you can modify the name.
 - **Map replica virtual network IP to target networks**—You will need to identify how you want to handle the network mapping. The **Replica VM Network IP** lists the IP addresses contained in the image of the original source. Map each one to a **Target Network**, which is a virtual network on the ESX server.
 - **Number of processors**—Specify how many processors to configure on the virtual machine. The number of processors from the original source image is displayed.
 - **Amount of memory**—Specify the amount of memory to configure on the virtual machine. The amount of memory from the original source image is displayed.
15. Click **Next** to continue.
 16. Specify your recovery shutdown options.

Set recovery shutdown options



Manual or Automatic Recovery

The recovery process can continue automatically after a mirror completes, or it can replicate changes while waiting for user intervention.

Wait for user intervention before recovery



Shutdown Source Server


If the source server is detected on the network, it can be shutdown before the final reboot of the server to avoid identity conflicts.

Shut down the source server



Original Protection Job

Any updates to the image during recovery could corrupt the recovered server.

 If the original source server is already unavailable, this choice will not be observed.

Delete the original protection job.

Stop the original protection job. The job can be restarted after recovery.

Leave the original protection job running. The recovered server could be corrupted, however the original source will remain protected.

-
- **Wait for user intervention before recovery**—Specify if you want to pause the recovery process after the original source image has been mirrored to the Virtual Recovery Appliance (the existing virtual machine), which allows you time to coordinate any other tasks. If you do not pause the recovery, after the mirror the Virtual Recovery Appliance will immediately unmount the disk(s), remove them, and start the new virtual machine.
 - **Shut down the source server**—Specify if you want to shutdown the source server before the new virtual machine is started, which avoids identity conflicts on the network.
 - **Original Protection Job**—Specify what to do with the original protection job if the original source is still online. If the original source is not running, these options will be discarded because they cannot be performed.
 - **Delete the original protection job**—The original protection job is stopped and then deleted.
 - **Stop the original protection job**—The original protection job is stopped but is not deleted. It can be restarted after the recovery.
 - **Leave the original protection job running**—The original protection job is not stopped or deleted. Keep in mind with this option that the data on the recovery server could become inconsistent or corrupted because the original source protection job could be transmitting data to the repository server and that data might get transmitted to the recovery server, depending on the progress of the recovery process .
17. Click **Next** to continue.
 18. At the **Recovery summary** page, you can review the choices you have made and you can also set optional data transmission and e-mail notification settings. If you need to make any changes to the selections you made, click **Back** to return to the previous pages. If you want to [configure optional settings](#), click **Configure**. Once you have completed your configuration, click **Finish** to start the recovery.
 19. Monitor the progress of the recovery mirror on the [Monitor](#) page.
 20. When the recovery mirror has completed, the new virtual machine on the recovery server will automatically start to complete the recovery process. If you specified to pause the recovery, when the mirror is complete, the **Activity** will change to **Protected**. Use this time to complete any necessary tasks. When you are ready to complete the recovery, click **Recover** on the **Monitor** page toolbar to complete the recovery process. Once the virtual machine has been started, it will become the original source server.

Note Because the Windows product activation is dependent on hardware, you may need to reactivate your Windows registration after recovery. Follow the on-screen prompts to complete the reactivation.

Recovering an entire server to an automatically provisioned virtual machine on Hyper-V

1. There are two ways to begin the recovery process.
 - Click **Recover** from the left navigation pane and identify the repository server that contains your source data that you want to recover. Click **Next** to continue.
 - Highlight the data protection job on the **Monitor** page and select **Recover** from the toolbar. This automatically identifies the repository server that contains the source data that you want to recover.
2. Select the image of the source that you want to recover and the point-in-time strategy.

Choose image



Source server image

Select an image

Available images:

Image Name	Image Location	Image State	Browse...
alpha	c:/server image/alpha	Good	



Point-in-time strategy

Live data

Snapshot

Date	Type	Status
7/23/2009 12:53:55 PM	User Request	Ok

Specific time

Data is valid for recovery during these times:

Start Time	End Time
7/23/2009 12:53:55 AM	7/25/2009 1:04:39 PM
7/25/2009 1:07:02 PM	7/31/2009 2:17:21 PM
7/31/2009 3:58:24 PM	8/12/2009 11:44:49 AM

- **Available images**—Select the image of the source that contains the data you want to recover. If the image is not displayed, click **Browse** to look for the location of the server image.
 - **Live data**—Select this option to recover to the live data that is currently stored on the repository server.
 - **Snapshot**—Select this option and choose a snapshot from the list to recover to the data stored in that snapshot. If you did not enable snapshots or there were no snapshots taken on the repository server, there will be no snapshots in the list.
 - **Specific time**—Select this option and choose a specific date and time to recover the data from that point in time. The table below the date and time fields allows you to confirm that the Double-Take Backup data is valid for the point in time you are selecting. If you did not enable continuous data protection when you configured protection or if the TimeData initialization has not completed, this option will not be available.
3. Click **Next** to continue.
 4. Select the recovery server. This is the Hyper-V server that will host the virtual machine that, after the recovery, will become your new source.

Choose recovery server



Choose recovery server

The source image can be recovered to an existing server, or a new virtual machine can be created for the recovery. To create a virtual machine for this recovery, select the option below and specify either a Microsoft Hyper-V or a VMware ESX host.

Provision a new virtual machine for this recovery.

Server:
gamma

User name:
administrator

Password:
●●●●●●●●

Domain:

Supplied credentials will be used continuously to access this server.
The domain entry will be ignored for a VMware ESX host.

- **Provision a new virtual machine for this recovery**—Enable this option to recover to an automatically provisioned virtual machine hosted on a Hyper-V server.
 - **Server**—Specify the name of the Hyper-V server. You can also click **Browse** to select the name from Active Directory or from a network drill-down list.
 - **User name**—Specify a user that is a member of the **Double-Take Admin** security group on the recovery server. If your original source was the only domain controller in your network, you must specify a local account in the local administrators group.
 - **Password**—Specify the password associated with the **User name** you entered.
 - **Domain**—If you are working in a domain environment, specify the **Domain**.
5. Click **Next** to continue.
 6. Because you are creating a new virtual machine, you must **Recover the entire server** so that the system state from the original source can be used to create the new virtual. You can exclude volumes that you protected but do not want to recover, but be careful when excluding data. Excluded volumes may compromise the integrity of your applications. Volumes marked with a checkmark are included. If there is no checkmark, the item is excluded.

Choose image data to recover



Source data

Choose the source data to recover. System state file and folder cannot be modified.

- Recover the entire server.
- Recover selected data only. This will not perform a full server recovery.

<input checked="" type="checkbox"/>	C:
<input checked="" type="checkbox"/>	E:
<input checked="" type="checkbox"/>	F:

7. Click **Next** to continue.
8. Specify the volumes to create on the new virtual machine.

Edit provisioned volumes



Provisioned Volumes

Edit the provisioned volumes on the virtual machine that will take over for the source server after recovery.

Volume	Selected Data (GB)	Source Size (GB)	Provisioned Volume (GB)	Controller Type
C	5.91	9.99	9.99	IDE
E	0.10	4.99	4.99	IDE
F	218.68	512.65	512.65	IDE

- **Selected Data**—This displays the amount of data, in gigabytes, currently on the source volume.
 - **Source Size**—This displays the size, in gigabytes, of the volume on the source.
 - **Provisioned Volume**—Specify the size, in gigabytes, of the volume to be created on the new virtual.
 - **Controller Type**—Specify the type of controller, IDE or SCSI, to be used for each volume. The boot volume must be an IDE controller.
8. Click **Next** to continue.
 9. Specify the volume and, if desired a folder, on the target server where the new virtual machine will be created. You can only select a volume that has enough free space available.

Choose Replica Virtual Machine Home Folder



Choose Replica Virtual Machine Home Folder

Choose the volume and folder on the target server that will hold the replica virtual machine. Only volumes with at least 527.63 GB free space can be selected.

	Volume	Total Size (GB)	Free Space (GB)
	C	858.59	637.24
	E	5351.51	4496.75

10. Click **Next** to continue.
11. Configure the new virtual machine that will be created and will become the source after the recovery.

Edit replica settings



Replica virtual machine

Configure the virtual machine that will take over for the source server after cutover.

Hyper-V Manager display name

alpha_VM

Map replica virtual network IP to target networks:

Replica VM Network IP	Target Network
172.31.246.166	Broadcom BCM5708C NetXtreme II GigE (NDIS VBD Client) #2 - Virtual N...

Number of processors:

1

Processors on the source server:

1

Amount of memory (MB):

512

Memory on the source server (MB):

512

- **Hyper-V Manager display name**—By default, source name_VM will be used to create the new virtual machine display name. If desired, you can modify the name.
 - **Map replica virtual network IP to target networks**—You will need to identify how you want to handle the network mapping. The **Replica VM Network IP** lists the IP addresses contained in the image of the original source. Map each one to a **Target Network**, which is a virtual network on the Hyper-V server.
 - **Number of processors**—Specify how many processors to configure on the virtual machine. The number of processors from the original source image is displayed.
 - **Amount of memory**—Specify the amount of memory to configure on the virtual machine. The amount of memory from the original source image is displayed.
12. Click **Next** to continue.
 13. Specify your recovery shutdown options.

Set recovery shutdown options



Manual or Automatic Recovery

The recovery process can continue automatically after a mirror completes, or it can replicate changes while waiting for user intervention.

Wait for user intervention before recovery



Shutdown Source Server

If the source server is detected on the network, it can be shutdown before the final reboot of the server to avoid identity conflicts.

Shut down the source server



Original Protection Job

Any updates to the image during recovery could corrupt the recovered server.

If the original source server is already unavailable, this choice will not be observed.

Delete the original protection job.

Stop the original protection job. The job can be restarted after recovery.

Leave the original protection job running. The recovered server could be corrupted, however the original source will remain protected.

- **Wait for user intervention before recovery**—Specify if you want to pause the recovery process after the original source image has been mirrored to the new virtual machine on the Hyper-V server, which allows you time to coordinate any other tasks. If you do not pause the recovery, the new virtual machine will be started immediately after the mirror is complete.

-
- **Shut down the source server**—Specify if you want to shutdown the source server before the new virtual machine is started, which avoids identity conflicts on the network.
 - **Original Protection Job**—Specify what to do with the original protection job if the original source is still online. If the original source is not running, these options will be discarded because they cannot be performed.
 - **Delete the original protection job**—The original protection job is stopped and then deleted.
 - **Stop the original protection job**—The original protection job is stopped but is not deleted. It can be restarted after the recovery.
 - **Leave the original protection job running**—The original protection job is not stopped or deleted. Keep in mind with this option that the data on the recovery server could become inconsistent or corrupted because the original source protection job could be transmitting data to the repository server and that data might get transmitted to the recovery server, depending on the progress of the recovery process .
14. Click **Next** to continue.
 15. At the **Recovery summary** page, you can review the choices you have made and you can also set optional data transmission and e-mail notification settings. If you need to make any changes to the selections you made, click **Back** to return to the previous pages. If you want to [configure optional settings](#), click **Configure**. Once you have completed your configuration, click **Finish** to start the recovery.
 16. Monitor the progress of the recovery mirror on the [Monitor](#) page.
 17. When the recovery mirror has completed, the new virtual machine on the recovery server will automatically start to complete the recovery process. If you specified to pause the recovery, when the mirror is complete, the **Activity** will change to **Protected**. Use this time to complete any necessary tasks. When you are ready to complete the recovery, click **Recover** on the **Monitor** page toolbar to complete the recovery process. Once the virtual machine has been started, it will become the original source server.

Note Because the Windows product activation is dependent on hardware, you may need to reactivate your Windows registration after recovery. Follow the on-screen prompts to complete the reactivation.

Recovering data to an existing physical or virtual machine

1. There are two ways to begin the recovery process.
 - Click **Recover** from the left navigation pane and identify the repository server that contains your source data that you want to recover. Click **Next** to continue.
 - Highlight the data protection job on the **Monitor** page and select **Recover** from the toolbar. This automatically identifies the repository server that contains the source data that you want to recover.
2. Select the image of the source that you want to recover and the point-in-time strategy.

Choose image



Source server image

Select an image

Available images:

Image Name	Image Location	Image State	Browse...
alpha	c:/server image/alpha	Good	



Point-in-time strategy

Live data

Snapshot

Date	Type	Status
7/23/2009 12:53:55 PM	User Request	Ok

Specific time

Data is valid for recovery during these times:

Start Time	End Time
7/23/2009 12:53:55 AM	7/25/2009 1:04:39 PM
7/25/2009 1:07:02 PM	7/31/2009 2:17:21 PM
7/31/2009 3:58:24 PM	8/12/2009 11:44:49 AM

- **Available images**—Select the image of the source that contains the data you want to recover. If the image is not displayed, click **Browse** to look for the location of the server image.
 - **Live data**—Select this option to recover to the live data that is currently stored on the repository server.
 - **Snapshot**—Select this option and choose a snapshot from the list to recover to the data stored in that snapshot. If you did not enable snapshots or there were no snapshots taken on the repository server, there will be no snapshots in the list.
 - **Specific time**—Select this option and choose a specific date and time to recover the data from that point in time. The table below the date and time fields allows you to confirm that the Double-Take Backup data is valid for the point in time you are selecting. If you did not enable continuous data protection when you configured protection or if the TimeData initialization has not completed, this option will not be available.
3. Click **Next** to continue.
 4. You will be prompted that you are restoring data only and not system state information. Continue the recovery by selecting **Yes**.
 5. Select the recovery server. If you used the Double-Take Backup built-in installation to push a copy of Double-Take Backup to the recovery server, that copy will automatically be activated when you select it as a recovery server. The copy will automatically deactivate after five days.

Choose recovery server



Choose recovery server

The source image can be recovered to an existing server, or a new virtual machine can be created for the recovery. To create a virtual machine for this recovery, select the option below and specify either a Microsoft Hyper-V or a VMware ESX host.

Provision a new virtual machine for this recovery.

Server:
gamma

User name:
administrator

Password:
●●●●●●●●

Domain:
domain.com

Supplied credentials will be used continuously to access this server.

- **Provision a new virtual machine for this recovery**—This option will not be available. It is only available for full-server recoveries.
 - **Server**—Specify the name of the recovery server. You can also click **Browse** to select the name from Active Directory or from a network drill-down list.
 - **User name**—Specify a user that is a member of the **Double-Take Admin** security group on the recovery server.
 - **Password**—Specify the password associated with the **User name** you entered.
 - **Domain**—If you are working in a domain environment, specify the **Domain**.
6. Click **Next** to continue.
7. Because you are working with data only, you will be unable to **Recover the entire server**. You can only recover the data you selected to protect. You can exclude files that you protected but do not want to recover, but be careful when excluding data. Excluded volumes, folders, and/or files may compromise the integrity of your applications.

Choose image data to recover



Source data

Choose the source data to recover. System state file and folder cannot be modified.

- Recover the entire server.
- Recover selected data only. This will not perform a full server recovery.

<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	C:
<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	E:
<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	F:

Volumes, folders, and files that are marked with a checkmark are included. If there is no checkmark, the item is excluded. Expand and collapse the directory tree and click on an item to add or remove a checkmark. Once you have configured the data to recover, click **Next** to continue.

Note:

The **Select all** and **Deselect all** buttons allow you to quickly select or deselect all data on your server based on your system state or selected data only choice. These buttons will do not change what is selected for the system state, which is defined by Double-Take Backup and cannot be modified.

- Specify the location, including the volume, on the recovery server to store the source data. For example, you might enter C: or C:\directory.
- Click **Next** to continue.
- Specify your recovery shutdown options.

Set recovery shutdown options



Manual or Automatic Recovery

The recovery process can continue automatically after a mirror completes, or it can replicate changes while waiting for user intervention.

- Wait for user intervention before recovery



Shutdown Source Server

If the source server is detected on the network, it can be shutdown before the final reboot of the server to avoid identity conflicts.

- Shut down the source server



Original Protection Job

Any updates to the image during recovery could corrupt the recovered server.

- If the original source server is already unavailable, this choice will not be observed.
- Delete the original protection job.
- Stop the original protection job. The job can be restarted after recovery.
- Leave the original protection job running. The recovered server could be corrupted, however the original source will remain protected.

- **Wait for user intervention before recovery**—This option is not available for data recovery. It is for full-server recovery jobs only.
 - **Shut down the source server from the original protection job**—Specify if you want to shutdown the source server after the source data is mirrored to the recovery server. This option is not available if the recovery server is the original source.
 - **Original Protection Job**—Specify what to do with the original protection job if the original source is still online. If the original source is not running, these options will be discarded because they cannot be performed. This option is not available if the recovery server is the original source or the repository server.
 - **Delete the original protection job**—The original protection job is stopped and then deleted.
 - **Stop the original protection job**—The original protection job is stopped but is not deleted. It can be restarted after the recovery.
 - **Leave the original protection job running**—The original protection job is not stopped or deleted. Keep in mind with this option that the data on the recovery server could become inconsistent or corrupted because the original source protection job could be transmitting data to the repository server and that data might get transmitted to the recovery server, depending on the progress of the recovery process .
- Click **Next** to continue.
 - At the **Recovery summary** page, you can review the choices you have made and you can also set optional data transmission and e-mail notification settings. If you need to make any changes to the selections you made, click **Back** to return to the previous pages. If you want to [configure optional settings](#), click **Configure**. Once you have completed your configuration, click **Finish** to start the recovery.
 - Monitor the progress of the recovery mirror on the [Monitor](#) page.

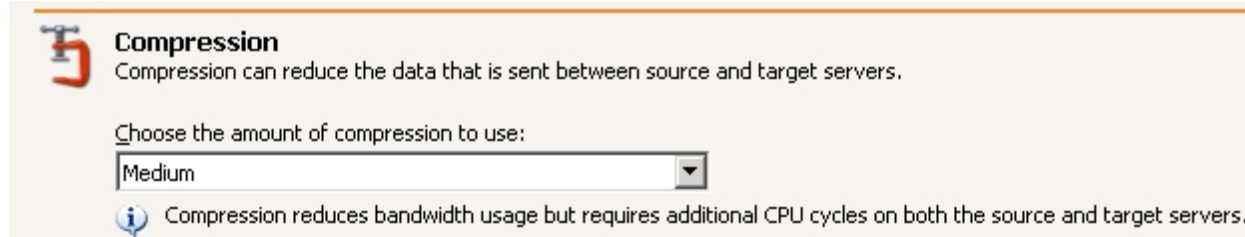
Optional recovery settings

When you are creating a recovery job, you will have an opportunity to configure optional job settings when you get to the **Recovery summary** workflow page. If you want to configure optional settings, click **Configure** next to the options section.

-
- [Compressing data for transmission for a recovery job](#)
 - [Limiting bandwidth for a recovery job](#)
 - [Routing transmission for a recovery job](#)
 - [E-mailing recovery server event messages](#)

Compressing data for transmission for a recovery job

To help reduce the amount of bandwidth needed to transmit Double-Take Backup 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 repository server receives the compressed data, it decompresses it and then writes it to disk.

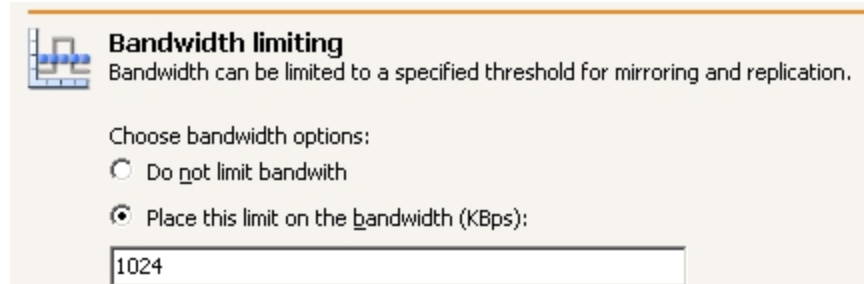


The screenshot shows a panel titled "Compression" with a red icon of a hand holding a pencil. Below the title, it says "Compression can reduce the data that is sent between source and target servers." There is a dropdown menu labeled "Choose the amount of compression to use:" with "Medium" selected. Below the dropdown is an information icon and the text "Compression reduces bandwidth usage but requires additional CPU cycles on both the source and target servers."

By default, compression is set to **None**, which means it is disabled. The compression levels can be set at **Minimum**, **Medium**, or **Maximum**. Each level requires more overhead than the previous level of compression. 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.

Limiting bandwidth for a recovery job

Bandwidth limitations are available to restrict the amount of network bandwidth used for Double-Take Backup data transmissions. When a bandwidth limit is specified, Double-Take Backup never exceeds that allotted amount. The bandwidth not in use by Double-Take Backup is available for all other network traffic.

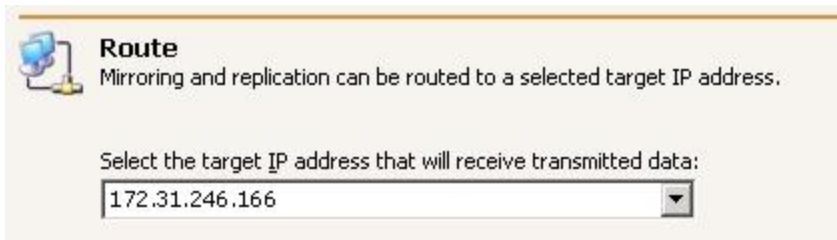


The screenshot shows a panel titled "Bandwidth limiting" with a blue icon of a network diagram. Below the title, it says "Bandwidth can be limited to a specified threshold for mirroring and replication." There are two radio button options: "Do not limit bandwidth" (unselected) and "Place this limit on the bandwidth (KBps):" (selected). Below the selected option is a text input field containing the value "1024".

- **Do not limit bandwidth**—Double-Take Backup will transmit data using 100% bandwidth availability.
- **Place this limit on the bandwidth**—Enter a value in kilobytes per second to limit data transmission. This is the maximum amount of data that will be transmitted per second.

Routing transmission for a recovery job

By default, Double-Take Backup will select a default route for transmissions. The default route is the first connection listed in the NIC binding order on the repository server. If you want to select a different route for Double-Take Backup transmissions, you can select a different IP address on the repository server using the Double-Take Backup **Route** option, or you can reorder the bindings on the repository server (**Network Connections, Advanced, Advanced Settings, Adapters and Bindings**).

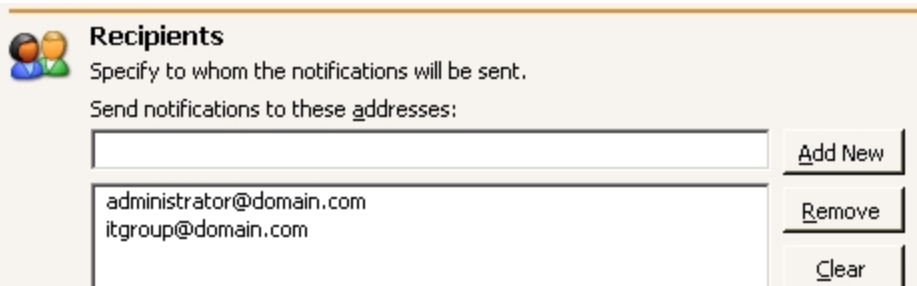


E-mailing recovery server event messages

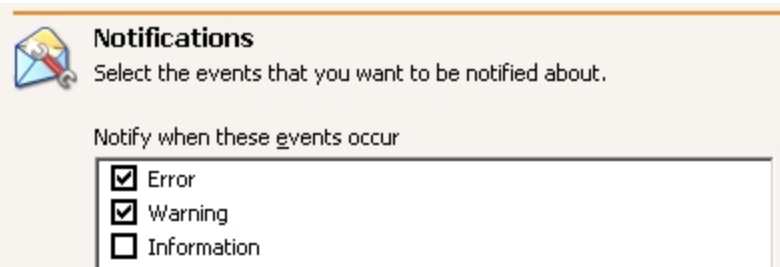
E-mail notification sends specific Double-Take Backup events to user-defined e-mail addresses. If you want to use e-mail notification, you will need to identify who will receive the messages, what messages they will receive, and your e-mail server configuration.

The e-mail notification settings are for the recovery server. If you want to configure e-mail notification for your repository server (which is acting in a source server role), see [E-mailing source event messages](#).


- **Recipients**—Specify the e-mail address that each Double-Take Backup e-mail message should be sent to and click **Add New**. The e-mail address will be inserted into the list of addresses. Each address is limited to 256 characters, and 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**. If you want to remove all of the addresses from the list, click **Clear**.



- **Notifications**—Specify which events that you want to be sent via e-mail. You can select **Error**, **Warning**, and/or **Information**.



- **Server Setup**—Specify your e-mail server configuration.



Server Setup

Configure from where and whom e-mail notifications will originate.

Mail Server (SMTP):
SMTPServerName Log on to SMTP server

User name:
administrator

Password:

From Address:
administrator@domain.com

Subject prefix:
Double-Take Software Notification Add event description to prefix

- **Mail Server**—Specify the name of your SMTP mail server.
- **Log on to SMTP server**—If your SMTP server requires a login, enable this option and specify a **User name** and **Password**.

Note: TimeData events do not support authenticated SMTP notification methods. Check the [TimeData log files](#) for TimeData events if you are using authenticated SMTP.

- **User name**—Specify a user account with privileges to send e-mail messages from your SMTP server.
- **Password**—Specify the password associated with the **User Name** you entered.
- **From Address**—Specify the e-mail address that you want to appear in the **From** field of each Double-Take Backup e-mail message. The address is limited to 256 characters.
- **Subject prefix**—Specify default text that will appear on every Double-Take Backup e-mail message. This will help distinguish Double-Take Backup messages from other messages.
- **Add event description to prefix**—Specify if you want the description of the event appended to the **Subject prefix**. The subject line is limited to 150 characters, so only the first 150 characters of the combined **Subject prefix** and event description will be displayed.

Note: Make sure you configure your e-mail client so that the Double-Take Backup e-mail notification messages do not get blocked by a spam filter or sent to a junk mail folder.

Job monitoring

Once a job is established you will want to monitor it. You can monitor the job directly from the Double-Take Backup console, or you can use several general monitoring tools that are available.

- [Monitoring a job through the Double-Take Backup console](#)
- [Monitoring log files](#)
- [Monitoring event messages](#)
- [Statistics](#)
- [Performance Monitor](#)
- [SNMP](#)
- [Error codes](#)

Monitoring a job through the Double-Take Backup console

Once you have a job established or imported, you can easily monitor the status of the job from the **Monitor** page.

When viewing the **Monitor** page, the main window is divided into two sections. The top pane displays overview monitoring information, while the bottom pane displays detailed data. You can hide the details pane by clicking on the **Hide details** text. You can display the details pane by clicking on the **Show details** text.

There is also a toolbar on the **Monitor** page for specific job and monitoring controls.

Job Type	Status	Source	Target	Time Remaining	% Complete
Server Protection	Protecting	alpha	beta	119 Minutes	16.3%





Hide details

Job properties		Connection properties	
Job type:	Server Protection	Total sent data:	149.80 MB
Job Status:	Protecting	Total sent compressed:	149.80 MB
Time Remaining:	119 Minutes	Transmit status:	Transmitting
Source server:	alpha	Target status:	Processing
Target server:	beta	Recovery available:	No
Image Location:	c:/server image/alpha		

Mirror statistics		Replication statistics	
Mirror status:	Mirroring	Replication status:	Replicating
Mirror percent complete:	16.3%	Bytes sent:	5.61 MB
Mirror remaining:	10.19 GB	Bytes sent compressed:	5.61 MB
Mirror skipped:	0 bytes	Bytes queued:	0 bytes

Overview status information

Overview information is provided for each job at the top of the **Monitor** page.

Column	Description
<no label>	<p>The first unlabeled column is an icon that indicates at-a-glance the status of the job.</p> <ul style="list-style-type: none">  The black, moving circular icon indicates the console is attempting to process the job to determine the status and statistics.  The white checkmark inside a green circle indicates the job is in a good state.  The black exclamation point inside a yellow triangle indicates the job is in a warning state.  The white X inside a red circle indicates the job is in a bad state.
Job Type	<ul style="list-style-type: none"> • Data Protection—This type of job is protecting selected data on the source. • Server Protection—This type of job is protecting an entire source, including the system state. • Data Recovery—This type of job is recovering selected data from the repository server • Server Recovery—This type of job is recovering an entire source from the repository server.
Status	<p>There are many different status messages that keep you informed of the progress of your jobs. Most of the status messages are informational and do not require any administrator interaction. If you see error messages, check the job details. You can also view the Double-Take log file on the server for additional information about processing on that machine.</p>
Source	<p>Source indicates which server is transmitting mirror and replication data.</p> <ul style="list-style-type: none"> • Data Protection jobs—When the Job Type is Data Protection, this is the name of your source. The data from the source is being transmitted from the source and is being received by the repository server. • Server Protection jobs—When the Job Type is Server Protection, this is the name of your source. The image of the source is being transmitted from the source and is being received by the repository server. • Data Recovery jobs—When the Job Type is Data Recovery, this is the name of the repository server. The data from the original source is being transmitted from the repository server and is being received by the recovery server. • Sever Recovery jobs—When the Job Type is Server Recovery, this is the name of the repository server. The image of the original source is being transmitted from the repository server and is being received by the recovery server.

Column	Description
Target	<p>Target indicates which server is receiving mirror and replication data.</p> <ul style="list-style-type: none"> • Data Protection jobs—When the Job Type is Data Protection, this is the name of your repository server. The data from the source is being transmitted from the source and is being received by the repository server. • Server Protection jobs—When the Job Type is Server Protection, this is the name of your repository server. The image of the source is being transmitted from the source and is being received by the repository server. • Data Recovery jobs—When the Job Type is Data Recovery, this is the name of the recovery server. The data from the original source is being transmitted from the repository server and is being received by the recovery server. • Server Recovery jobs—When the Job Type is Server Recovery, this is the name of the recovery server. The image of the original source is being transmitted from the repository server and is being received by the recovery server.
Time Remaining	When mirroring, this is the estimated time remaining to complete the mirror. For various reasons, this calculation is an estimate. For example, some file components, like alternate data streams, cannot be included in mirror calculations, or the size of the data set may have changed after the mirror started.
% Complete	When mirroring, this is the estimated percentage of the mirror that is complete. For various reasons, this calculation is an estimate. For example, some file components, like alternate data streams, cannot be included in mirror calculations, or the size of the data set may have changed after the mirror started.

Detailed status information

Detailed information is provided for each job at the bottom of the **Monitor** page.

Section	Detailed Item	Description
Job properties	Job type	<ul style="list-style-type: none"> • Data Protection—This type of job is protecting selected data on the source. • Server Protection—This type of job is protecting an entire source, including the system state. • Data Recovery—This type of job is recovering selected data from the repository server • Server Recovery—This type of job is recovering an entire source from the repository server.












Section	Detailed Item	Description
Job properties	Job Status	<p>There are many different status messages that keep you informed of the progress of your jobs. Most of the status messages are informational and do not require any administrator interaction. If you see error messages, check the job details.</p> <p>If the Job Status is Communication Error, the issue could be several things such as the source server is unavailable or the client has invalid credentials for the source. If it is a problem with the source server, perhaps the Double-Take service is not running, resolve the problem on the source. If the source appears to have no issues, try refreshing the job credentials through the Update credentials toolbar button.</p> <p>For more details on error messages, review the log file on the server for additional information about processing on that machine.</p>
Job properties	Time Remaining	<p>When mirroring, this is the estimated time remaining to complete the mirror. For various reasons, this calculation is an estimate. For example, some file components, like alternate data streams, cannot be included in mirror calculations, or the size of the data set may have changed after the mirror started.</p>
Job properties	Source server	<p>Source server indicates which server is transmitting mirror and replication data.</p> <ul style="list-style-type: none"> • Data Protection jobs—When the Job Type is Data Protection, this is the name of your source. The data from the source is being transmitted from the source and is being received by the repository server. • Server Protection jobs—When the Job Type is Server Protection, this is the name of your source. The image of the source is being transmitted from the source and is being received by the repository server. • Data Recovery jobs—When the Job Type is Data Recovery, this is the name of the repository server. The data from the original source is being transmitted from the repository server and is being received by the recovery server. • Sever Recovery jobs—When the Job Type is Server Recovery, this is the name of the repository server. The image of the original source is being transmitted from the repository server and is being received by the recovery server.



Section	Detailed Item	Description
Job properties	Target server	<p>Target server indicates which server is receiving mirror and replication data.</p> <ul style="list-style-type: none"> • Data Protection jobs—When the Job Type is Data Protection, this is the name of your repository server. The data from the source is being transmitted from the source and is being received by the repository server. • Server Protection jobs—When the Job Type is Server Protection, this is the name of your repository server. The image of the source is being transmitted from the source and is being received by the repository server. • Data Recovery jobs—When the Job Type is Data Recovery, this is the name of the recovery server. The data from the original source is being transmitted from the repository server and is being received by the recovery server. • Sever Recovery jobs—When the Job Type is Server Recovery, this is the name of the recovery server. The image of the original source is being transmitted from the repository server and is being received by the recovery server.
Job properties	Image Location	This is the location on the Target server where the data or image is stored.
Connection Properties	Total sent data	This is the total amount of data (mirroring and replication) that has been sent from the Source server to the Target server .
Connection Properties	Total sent compressed	This is the total amount of compressed data (mirroring and replication) that has been sent from the Source server to the Target server . If compression is disabled, this statistic will be the same as Total sent data .
Connection Properties	Transmit status	<ul style="list-style-type: none"> • Transmitting—The job is transmitting data. • Stopped—The job is stopped. • Paused—The job is paused. • Error—There is a transmission error.
Connection Properties	Target status	<ul style="list-style-type: none"> • Processing—The Target server is processing data. • Paused—The Target server is paused.
Connection Properties	Recovery available	<ul style="list-style-type: none"> • Yes—The mirror is complete and replication is continuing. The data or image is valid for recovery. • No—The mirror is not yet complete, so the data or image is not yet valid for recovery.

Section	Detailed Item	Description
Mirror statistics	Mirror status	<ul style="list-style-type: none"> • Calculating—The amount of mirror data is being calculated. • Mirroring—Data is being mirrored. • Idle—There is no data to mirror. • Paused—The mirror is paused. • Stopped—The mirror is stopped. • Unknown—The console is trying or unable to determine the mirror status.
Mirror statistics	Mirror percent complete	This is the estimated percentage of the mirror that is complete. For various reasons, this calculation is an estimate. For example, some file components, like alternate data streams, cannot be included in mirror calculations, or the size of the data set may have changed after the mirror started.
Mirror statistics	Mirror remaining	This is the estimated amount of mirroring data that still has to be transmitted. For various reasons, this calculation is an estimate. For example, some file components, like alternate data streams, cannot be included in mirror calculations, or the size of the data set may have changed after the mirror started.
Mirror statistics	Mirror skipped	This is the amount of mirroring data that has been skipped (because the data is not different on the Source server and Target server).
Replication statistics	Replication status	<ul style="list-style-type: none"> • Replicating—Data is being replicated. • Pending—Replication is pending. • Stopped—Replication is stopped. • Out of memory—Kernel memory has been exhausted. • Watchdog—The Double-Take service is not receiving replication operations from the Double-Take driver. • Unknown—The console is trying or unable to determine the replication status.
Replication statistics	Bytes sent	This is the total amount of replication data that has been sent from the Source server to the Target server .
Replication statistics	Bytes sent compressed	This is the total amount of compressed replication data that has been sent from the Source server to the Target server . If compression is disabled, this statistic will be the same as Bytes sent .
Replication statistics	Bytes queued	This is the total amount of replication data that is in queue on the source.

Monitor page toolbar

The **Monitor** page has its own toolbar buttons for job and monitoring controls.

Toolbar Icon	Tooltip Text	Description
	New activity	The large shield icon on the far left of the toolbar opens a small menu that lets you select an action to perform. Select Protect or Recover .
	Import jobs	The smaller shield icon with the green circle with white plus sign allows you to launch the import jobs workflow.
	Stop monitoring job	The smaller shield icon with the red circle with white line removes the job from the console Monitor page. The job remains active, but is no longer visible from the console. To see the job again, import it using the Import jobs button. If you want to permanently delete the job, use the Delete job button.
	Resume	The green triangle icon resumes the selected job.
	Stop	The red square icon stops the selected job. The image of the source, at the time the job was stopped, is still available for recovery, but the image is no longer being updated with data changes from the source. If you want to use the job again, you must restart it using the Resume button. When you restart a job, a remirror will automatically be performed.
	Take a snapshot	The blue icon resembling a photograph takes an immediate snapshot of the data on the repository server. A pop-up dialog box will notify you of the success or failure of the snapshot.
	Recover	The life preserver icon starts the recovery process for the selected job.
	Update credentials	The keys icon allows you to specify credentials for the Source server and the Target server . If the Password associated with a specified User name changes after a job has been started, Double-Take Backup will attempt to update the credentials it is using. If the automatic update is not successful, the monitor page will report Invalid credentials . In this case, you will need to manually update the credentials that the job is using the Update credentials button.
	Delete job	The trash can icon permanently deletes the selected job. The image of the source, at the time the job was deleted, is still available for recovery, but the image is no longer being updated with data changes from the source. If you want to use the job again, you must re-create it using the New activity button.
	View details	The paper with magnifying glass icon toggles between displaying and hiding the details at the bottom of the Monitor page.
	OK	The green circle with the white checkmark icon toggles between displaying and hiding those jobs that are in a good state. The number of jobs in the good state is also displayed on the toolbar button.

Toolbar Icon	Tooltip Text	Description
	Warnings	The yellow triangle with the black exclamation point icon toggles between displaying and hiding those jobs that are in a warning state. The number of jobs in the warning state is also displayed on the toolbar button.
	Error	The red circle with the white X icon toggles between displaying and hiding those jobs that are in an error state. The number of jobs in the error state is also displayed on the toolbar button.
		To the right of all of the toolbar buttons is a summary of the number of jobs displayed by the good, warning, and error state toggle buttons. The total number of jobs available for display in this console is also displayed.

Monitoring log files

Various Double-Take Backup components, for example the Double-Take service and the Double-Take Backup console, generate a log file to gather alerts, which are notification, warning, and error messages. The log files are written to disk.

Log File	Description
livewire.log	This log records basic processing from the Double-Take Backup console. It can be found in the Double-Take Backup installation directory where the console is being run. This log contains informational messages only.
livewire_debug.log	This log records detailed processing from the Double-Take Backup console. It can be found in the Double-Take Backup installation directory where the console is being run.
jobmgr.log	This log records detailed processing from the Double-Take Backup servers. It can be found in the Double-Take Backup installation directory on the repository or recovery server.
JobName.log	This log records virtual server provisioning processing when recovering to provisioned virtual machines. The file name is based on the name of the job. It can be found in the \VRA\logs directory under the Double-Take Backup installation directory on the repository or recovery server.
dtlogX.dtl	This log contains information from the Double-Take service. The letter X in the file name represents a series number, for example dtlog1.dtl, dtlog2.dtl, dtlog3.dtl, and so on. It can be found in the Double-Take Backup installation directory on the Double-Take Backup servers.
TDAgent-Default.log TDMsgSvr.log TDRepSvr.log	These three log files record detailed processing for TimeData continuous data protection. They can be found in the \TimeData\logs directory under the Double-Take Backup installation directory on the repository.

Monitoring event messages

An event is a significant occurrence in the system or in an application that requires administrators to be notified. The operating system writes notifications for these events to a log that can be displayed using the Windows Event Viewer. Three different log files are generated: application, security, and system.

1. To access the Event Viewer, select **Programs, Administrative Tools, Event Viewer**.
2. Select the log to view (**System, Security, or Application**) from the left pane of the Event Viewer. The following information is displayed for an event in the right pane of the Event Viewer.
 - **Type**—A classification of the event, such as Error, Warning, Information, Success Audit, or Failure Audit.
 - **Date**—The date the event occurred.
 - **Time**—The time the event occurred.
 - **Source**—The software that logged the event, which can be either an application or a component of the system, such as a driver.
 - **Category**—A classification of the event.
 - **Event**—Shows an ID number to identify the specific event. The **Event** helps product-support representatives track events in the system.
 - **User**—Identifies the user that logged the event.
 - **Computer**—The name of the computer where the event occurred.
3. To view a detailed description, double-click an event. The additional information is displayed in the Event Properties screen.

For additional information on customizing the Event Viewer (such as sorting the display, filtering the display, and so on), see your Windows reference guide or the Windows online help.

Event messages

The following table identifies the events.

ID	Category	Severity	Event Message	Required Response
1	Activation Key	Error	This evaluation period has expired. Mirroring and replication have been stopped. To obtain a license, please contact your vendor.	Contact your vendor to purchase either a single or site license.
2	Activation Key	Info.	The evaluation period expires in %1 day(s).	Contact your vendor before the evaluation period expires to purchase either a single or site license.

ID	Category	Severity	Event Message	Required Response
3	Activation Key	Info.	The evaluation period has been activated and expires in %1 day(s).	Contact your vendor before the evaluation period expires to purchase either a single or site license.
4	Activation Key	Warning	Duplicate activation codes detected on machine %1 from machine %2.	If you have an evaluation license or a site license, no action is necessary. If you have a single license, you must purchase either another single license or a site license.
5	Activation Key	Error	This product edition can only be run on Windows Server or Advanced Server running the Server Appliance Kit.	Verify your activation code has been entered correctly and contact technical support.
1000	DTCounters	Error	An exception occurred: %1	Run the installation and select Repair. Contact technical support if this event occurs again.
1001	DTCounters	Error	The Double-Take counter DLL could not initialize the statistics handler object to gather performance data.	Run the installation and select Repair. Contact technical support if this event occurs again.
1002	DTCounters	Error	The Double-Take counter DLL could not map shared memory file containing the performance data.	Run the installation and select Repair. Contact technical support if this event occurs again.
1003	DTCounters	Error	The Double-Take counter DLL could not open the "Performance" key in the Double-Take section of the registry.	Run the installation and select Repair. Contact technical support if this event occurs again.
1004	DTCounters	Error	The Double-Take counter DLL could not read the "First Counter" value under the Double-Take\Performance Key.	Run the installation and select Repair. Contact technical support if this event occurs again.
1005	DTCounters	Error	The Double-Take counter DLL read the "First Help" value under the Double-Take\Performance Key.	Run the installation and select Repair. Contact technical support if this event occurs again.

ID	Category	Severity	Event Message	Required Response
1006	DTCounters	Error	The Double-Take counter DLL could not create event handler for the worker thread.	Run the installation and select Repair. Contact technical support if this event occurs again.
3000	Service	Info.	Logger service was successfully started.	No action required.
3001	Service	Info.	Logger service was successfully stopped.	No action required.
4000	Service	Info.	Kernel was successfully started.	No action required.
4001	Service	Info.	Target service was successfully started.	No action required.
4002	Service	Info.	Source service was successfully started.	No action required.
4003	Service	Info.	Source service was successfully stopped.	No action required.
4004	Service	Info.	Target service was successfully stopped.	No action required.
4005	Service	Info.	Kernel was successfully stopped.	No action required.
4006	Service	Error	Service has aborted due to the following unrecoverable error: %1	Restart the Double-Take service.
4007	Service	Warning	Auto-disconnecting from %1 (%2) for Replication Set %3, ID: %4 due to %5	The connection is auto-disconnecting because the disk-based queue on the source has been filled, the 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 Backup cannot read from or write to the transaction log file).
4008	Service	Info.	Auto-disconnect has succeeded for %1 (%2) for Replication Set %3, ID: %4	No action required.

ID	Category	Severity	Event Message	Required Response
4009	Service	Info.	Auto-reconnecting Replication Set %1 to %2 (%3)	No action required.
4010	Service	Info.	Auto-reconnect has succeeded connecting Replication Set %1 to %2 (%3)	No action required.
4011	Service	Error	Auto-reconnect has failed connecting Replication Set %1 to %2 (%3)	Manually reestablish the replication set to target connection.
4012	Service	Warning	%1	This is a placeholder message for many other messages. See the specific log message for additional details.
4013	Service	Info.	%1	This is a placeholder message for many other messages. See the specific log message for additional details.
4014	Service	Info.	Service has started network transmission.	No action required.
4015	Service	Info.	Service has stopped network transmission.	No action required.
4016	Service	Info.	Service has established a connection to %1 (%2) for Replication Set %3, ID: %4	No action required.
4017	Service	Info.	Service has disconnected from %1 (%2) for Replication Set %3, ID: %4	No action required.
4018	Service	Warning	%1, however, mirroring and replication have been disabled as a restore is required due to a previous failover.	Perform a restoration.
4019	Service	Info.	Service has started a mirror to %1 (%2) for Replication Set %3, ID: %4	No action required.
4020	Service	Info.	Service has paused a mirror to %1 (%2) for Replication Set %3, ID: %4	No action required.

ID	Category	Severity	Event Message	Required Response
4021	Service	Info.	Service has resumed a mirror to %1 (%2) for Replication Set %3, ID: %4	No action required.
4022	Service	Info.	Service has stopped a mirror to %1 for Replication Set %2, ID: %3, %4	No action required.
4023	Service	Info.	Service has completed a mirror to %1 %2 for Replication Set %3, ID: %4, %5	No action required.
4024	Service	Info.	Service has started Replication to %1 (%2) for Replication Set %3, ID: %4	No action required.
4025	Service	Info.	Service has stopped Replication to %1 (%2) for Replication Set %3, ID: %4	No action required.
4026	Service	Info.	The target has been paused due to user intervention.	No action required.
4027	Service	Info.	The target has been resumed due to user intervention.	No action required.
4028	Service	Warning	Registration of service class with Active Directory failed. Verify that the Active Directory server is up and the service has the proper permissions to update its entries.	Verify that the Active Directory server is running and that the Double-Take service has permission to update Active Directory.
4029	Service	Warning	Registration of service instance with Active Directory failed. Verify that the Active Directory server is up and the service has the proper permissions to update its entries.	Verify that the Active Directory server is running and that the Double-Take service has permission to update Active Directory.
4030	Service	Error	RSResource.dll has an unknown error. The product functionality has been disabled.	Reinstall the software, using the installation Repair option, to install a new copy of the RSResource.dll. Contact technical support if this error persists.

ID	Category	Severity	Event Message	Required Response
4031	Service	Error	RSResource.dll could not be opened. The product functionality has been disabled.	Reinstall the software, using the installation Repair option, to install a new copy of the RSResource.dll. Contact technical support if this error persists.
4032	Service	Error	The RSResource.dll component version does not match the component version expected by the product. The product functionality has been disabled.	Reinstall the software, using the installation Repair option, to install a new copy of the RSResource.dll. Contact technical support if this error persists.
4033	Service	Error	RSResource.dll build version is invalid. The product functionality has been disabled.	Reinstall the software, using the installation Repair option, to install a new copy of the RSResource.dll. Contact technical support if this error persists.
4034	Service	Error	Error verifying the service name. The product functionality has been disabled.	Reinstall the software, using the installation Repair option, to install a new copy of the RSResource.dll. Contact technical support if this error persists.
4035	Service	Error	Error verifying the product name. The product functionality has been disabled.	Reinstall the software, using the installation Repair option, to install a new copy of the RSResource.dll. Contact technical support if this error persists.
4036	Service	Error	Error verifying the vendor name. The product functionality has been disabled.	Reinstall the software, using the installation Repair option, to install a new copy of the RSResource.dll. Contact technical support if this error persists.
4037	Service	Error	Error verifying the vendor URL name. The product functionality has been disabled.	Reinstall the software, using the installation Repair option, to install a new copy of the RSResource.dll. Contact technical support if this error persists.

ID	Category	Severity	Event Message	Required Response
4038	Service	Error	Error verifying the product code. The product functionality has been disabled.	Reinstall the software, using the installation Repair option, to install a new copy of the RSResource.dll. Contact technical support if this error persists.
4039	Service	Error	Error while reading RSResource.dll. The product functionality has been disabled.	Reinstall the software, using the installation Repair option, to install a new copy of the RSResource.dll. Contact technical support if this error persists.
4040	Service	Error	The product code is illegal for this computer hardware. The product functionality has been disabled.	Reinstall the software, using the installation Repair option, to install a new copy of the RSResource.dll. Contact technical support if this error persists.
4041	Service	Error	The product code is illegal for this operating system version. The product functionality has been disabled.	Reinstall the software, using the installation Repair option, to install a new copy of the RSResource.dll. Contact technical support if this error persists.
4042	Service	Error	The product code requires installing the Windows Server Appliance Kit. The product functionality has been disabled.	Reinstall the software, using the installation Repair option, to install a new copy of the RSResource.dll. Contact technical support if this error persists.
4043	Service	Error	This product can only be run on a limited number of processors and this server exceeds the limit. The product functionality has been disabled.	Reinstall the software, using the installation Repair option, to install a new copy of the RSResource.dll. Contact technical support if this error persists.
4044	Service	Error	An error was encountered and replication has been stopped. It is necessary to stop and restart the service to correct this error.	Contact technical support if this error persists.

ID	Category	Severity	Event Message	Required Response
4045	Service	Error	%1 value must be between 1025 and 65535. Using default of %2.	Verify that the Double-Take Backup port value you are trying to use is within the valid range. If it is not, it will automatically be reset to the default value.
4046	Service	Error	This service failed to start because of a possible port conflict. Win32 error: %1	Verify that the Double-Take Backup ports are not conflicting with ports used by other applications.
4047	Service	Error	Could not load ZLIB DLL %1. Some levels of compression will not be available.	The compression levels available depend on your operating system. You can reinstall the software, using the installation Repair option, to install a new copy of the DynaZip.dll, or contact technical support if this error persists.
4048	Service	Info.	Service has started a delete orphans task to %1 (%2) for Replication Set %3, ID: %4	No action required.
4049	Service	Info.	Service has paused a delete orphans task to %1 (%2) for Replication Set %3, ID: %4	No action required.
4050	Service	Info.	Service has resumed a delete orphans task to %1 (%2) for Replication Set %3, ID: %4	No action required.
4051	Service	Info.	Service has stopped a delete orphans task to %1 (%2) for Replication Set %3, ID: %4	No action required.
4052	Service	Info.	Service has completed a delete orphans task to %1 (%2) for Replication Set %3, ID: %4	No action required.
4053	Service	Info.	Service has started a restore task to %1 (%2) for Replication Set %3, ID: %4	No action required.
4054	Service	Info.	Service has paused a restore task to %1 (%2) for Replication Set %3, ID: %4	No action required.

ID	Category	Severity	Event Message	Required Response
4055	Service	Info.	Service has resumed a restore task to %1 (%2) for Replication Set %3, ID: %4	No action required.
4056	Service	Info.	Service has stopped a restore task to %1 (%2) for Replication Set %3, ID: %4	No action required.
4057	Service	Info.	Service has completed a restore task to %1 (%2) for Replication Set %3, ID: %4	No action required.
4058	Service	Info.	Service has started a verification task to %1 (%2) for Replication Set %3, ID: %4	No action required.
4059	Service	Info.	Service has paused a verification task to %1 (%2) for Replication Set %3, ID: %4	No action required.
4060	Service	Info.	Service has resumed a verification task to %1 (%2) for Replication Set %3, ID: %4	No action required.
4061	Service	Info.	Service has stopped a verification task to %1 (%2) for Replication Set %3, ID: %4	No action required.
4062	Service	Info.	Service has completed a verification task to %1 (%2) for Replication Set %3, ID: %4	No action required.
4063	Service	Info.	Bandwidth limit to %1 (%2) has changed to %3.	No action required.
4064	Service	Info.	Bandwidth limit to %1 (%2) is now in the "%3" period at %4.	No action required.
4065	Service	Warning	Target data state for connection %1 from %2 (%3) has changed because %4.	No action required.
4066	Service	Error	The product code requires a virtual server environment. The product functionality has been disabled.	The activation code you are using is for the Virtual Systems™ edition. This code will not work on non-virtual server environments.

ID	Category	Severity	Event Message	Required Response
4067	Service	Error	No replication ops have been received from the driver for an extended period of time.	Check other messages for errors with the Double-Take Backup drivers, and correct as required. If there are no driver messages, verify that your drives are connected to the source. If this error persists, contact technical support.
4068	Service	Error	Failed to write to a replicating volume.	Reboot the source server. Contact technical support if this event occurs again.
4096	System	Warning	The registry parameter %2 is unknown.	Delete the parameter and report this issue to technical support.
4097	System	Warning	Failed to initialize WMI support. The last Word in the Data Window is the NT status code.	No action required.
4097	System	Error	The file system filter failed to load. Replication will not occur. Reboot your server and contact technical support if this error occurs again. The last Word in the Data window is the NT status code.	Reboot your server and contact technical support if this event occurs again.
4098	System	Warning	The registry parameters failed to load, so the default configuration values will be used. The last Word in the Data window is the NT status code.	No action required.
4098	System	Error	The control device %2 was not created. Communication with the service will be disabled. Reboot the server and contact technical support if this error occurs again. The last Word in the Data window is the NT status code.	Reboot your server and contact technical support if this event occurs again.

ID	Category	Severity	Event Message	Required Response
4099	System	Warning	The driver detected a hard link for a file on drive %2. Hard links are not supported. Changes to this file will not be replicated.	Hard links are not supported.
4099	System	Error	The driver failed to register with filter manager. Reboot the server and contact technical support if this error occurs again. The last Word in the Data window is the NT status code.	Reboot your server and contact technical support if this event occurs again.
4100	Activation Key	Error	Product activation code is invalid. Please check that it is typed correctly and is valid for the version of the operating system in use.	If you are in the process of installing Double-Take Backup, verify that you are using a 24 character alpha-numeric code. If Double-Take Backup is already installed, confirm that the code entered is correct. If the code appears to be correct, contact technical support.
4100	System	Error	The versions of the driver and the filter driver do not match. Replication will not occur. Reboot your server. If this error occurs again, reinstall the software. Contact technical support if this error occurs after the software has been reinstalled. The last three Words in the Data window are the NT status code and the driver version numbers.	Reboot your server. Reinstall the software if this event occurs again. Contact technical support if this event occurs after reinstalling the software.
4101	Activation Key	Error	This service will not run on this device. Contact your sales representative for upgrade procedures.	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	Service	Warning	Target cannot write %1 due to target disk being full. Operation will be retried (%2 times or forever)	The disk on the target is full. The operation will be retried according to the TGExecutionRetryLimit setting.

ID	Category	Severity	Event Message	Required Response
4111	Service	Warning	Target can not write %1 due to a sharing violation. Operation will be retried (%2 times or forever)	A sharing violation error is prohibiting Double-Take Backup from writing on the target. The operation will be retried according to the TGExecutionRetryLimit setting.
4112	Service	Warning	Target can not write %1 due to access denied. Operation will be retried (%2 times or forever)	An access denied error is prohibiting Double-Take Backup from writing on the target. The operation will be retried according to the TGExecutionRetryLimit setting.
4113	Service	Warning	Target can not write %1 due to an unknown reason. Operation will be retried (%2 times or forever). Please check the log files for further information on the error.	An unknown error is prohibiting Double-Take Backup from writing on the target. The operation will be retried according to the TGExecutionRetryLimit setting.
4120	Service	Info.	Target write to %1 was completed successfully after %2 retries.	No action required.
4150	Service	Error	Target write %1 failed after %2 retries and will be discarded. See the event log or log files for error conditions. After correcting the problem, you should re-mirror or run a verify to resynchronize the changes.	The operation has been retried according to the TGExecutionRetryLimit setting but was not able to be written to the target and the operation was discarded. Correct the problem and remirror the files.
4155	Service	Warning	The service was unable to complete a file system operation in the allotted time. See the log files for error conditions. After correcting the problem, remirror or perform a verification with remirror to synchronize the changes.	Correct the file system error and then remirror or perform a verification with remirror to synchronize the changes.
4200	Service	Info.	In band task %1 submitted from %2 by %3 at %4	No action required.

ID	Category	Severity	Event Message	Required Response
4201	Service	Warning	In band task %1 discarded (submitted from %2 by %3 at %4)	A task may be discarded in the following scenarios: all connections to a target are manually disconnected, replication is stopped for all connections to a target, or an auto-disconnect occurs. If one of these scenarios did not cause the task to be discarded, contact technical support.
4202	Service	Info.	Running %1 in band script: %2 (task %3 submitted from %4 by %5 at %6)	No action required.
4203	Service	Info.	Completed run of in band script: %1 (exit code %2)	No action required.
4204	Service	Error	Error running in band script: %1	Review the task and its associated script(s) for syntax errors.
4205	Service	Warning	Timeout (%1 seconds) running in band script: %2	The timeout specified for the script to complete has expired. Normal processing will continue. You may need to manually terminate the script if it will never complete.
4206	Service	Warning	Run timeout disabled for in band script: %1	The timeout period was set to zero (0). Double-Take Backup will not wait for the script to complete before continuing. No action is required.
4207	Service	Warning	In band scripts disabled by server - no attempt will be made to run %1	Enable task command processing.
4300	Service	Error	A connection request was received on the target before the persistent target paths could be loaded.	You may need to disconnect and reconnect your replication set.
4301	Service	Error	Unable to block target paths, the driver is unavailable.	If you need to block your target paths, contact technical support.

ID	Category	Severity	Event Message	Required Response
4302	Service	Info.	Target Path %1 has been successfully blocked	No action required.
4303	Service	Warning	Blocking of target path: %1 failed. Error Code: %2	If you need to block your target paths, contact technical support.
4304	Service	Info.	Target Path %1 has been successfully unblocked	No action required.
4305	Service	Warning	Unblocking of target path: %1 failed. Error Code: %2	If you need to unblock your target paths, contact technical support.
4306	Service	Warning	Target paths for source %1 (%2) Connection id: %3 are already blocked	No action required.
4307	Service	Warning	Target paths for source %1 (%2) Connection id: %3 are already unblocked	No action required.
4308	Service	Error	Error loading target paths for blocking, registry key %1 has been corrupted.	If you need to block your target paths, contact technical support.
4400	Service	Error	Failed to create snapshot set for source %1 (%2) Connection ID: %3. Error: %4	The snapshot could not created. This may be due to a lack of disk space or memory or another reason. The error code is the Microsoft VSS error. Check your VSS documentation or contact technical support.
4401	Service	Error	Failed to delete automatic snapshot set for source %1 (%2) Connection ID: %3. Error: %4	The automatic snapshot could not deleted. This may be due to a lack of memory, the file does not exist, or another reason. The error code is the Microsoft Volume Shadow Copy error. Check your Volume Shadow Copy documentation or contact technical support.

ID	Category	Severity	Event Message	Required Response
4402	Service	Error	Failed to delete snapshot set for source %1 (%2) Connection ID: %3. Error: %4	The snapshot could not be deleted. This may be due to a lack of memory, the file does not exist, or another reason. The error code is the Microsoft Volume Shadow Copy error. Check your Volume Shadow Copy documentation or contact technical support.
4403	Service	Error	A scheduled snapshot could not be created for source %1 (%2) Connection ID: %3. because the target data was in a bad state. A snapshot will automatically be created when the target data reaches a good state.	No action required. A snapshot will automatically be created when the target data reaches a good state.
4404	Service	Info.	Set snapshot schedule for source %1 (%2) connection %3 to every %4 minutes. Next snapshot: %5.	No action required.
4405	Service	Info.	Removed snapshot schedule for source %1 (%2) connection %3.	No action required.
4406	Service	Info.	Enabled snapshot schedule for source %1 (%2) connection %3.	No action required.
4407	Service	Info.	Disabled snapshot schedule for source %1 (%2) connection %3.	No action required.
4408	Service	Warning	%1 was unable to move some orphans for source %2 on connection ID %3. Check the %1 logs for further details.	Orphan files could not be moved. For example, the location could be out of disk space. Check the Double-Take Backup log for more information.
4409	Service	Warning	%3 was unable to delete some orphans for source %1 on connection ID %2. Check the %3 logs for further details.	Orphan files could not be deleted. Check the Double-Take Backup log for more information.
4410	Service	Error	The registry hive dump failed with an of error: %1.	Contact technical support.

ID	Category	Severity	Event Message	Required Response
4411	Service	Warning	The Service has detected that port %1 is being %2 by the Windows Firewall.	The firewall port needs to be unblocked or restrictions against Double-Take Backup removed so that Double-Take Backup data can be transmitted.
5000	Service	Info.	Server Monitor service was successfully started.	No action required.
5001	Service	Info.	Server Monitor service was successfully stopped.	No action required.
5002	Service	Info.	Placeholders were modified to %1.	No action required.
5100	Failover	Info.	Failover completed for %1.	No action required.
5101	Failover	Info.	IP address %1 with subnet mask %2 was added to target machine's %3 adapter.	No action required.
5102	Failover	Warning	%1 has reached a failover condition. A response from the user is required before failover can take place.	User intervention has been configured. Open the Failover Control Center and accept or decline the failover prompt.
5103	Failover	Info.	Started adding drive shares from %1 to %2.	No action required.
5104	Failover	Info.	%1 drive shares were taken over by %2.	No action required.
5105	Failover	Info.	Attempting to run the %1 script.	No action required.
5106	Failover	Info.	The %1 script ran successfully.	No action required.
5107	Failover	Error	Error occurred in running %1 script.	Verify that the script identified exists with the proper permissions.
5108	Failover	Error	The source machine %1 is not responding to a ping.	This occurs when all monitored IP addresses on the source machine stop responding to pings. Countdown to failover will begin at the first occurrence and will continue until the source machine responds or until failover occurs.

ID	Category	Severity	Event Message	Required Response
5109	Failover	Error	The public NIC on source machine %1 is not responding to a ping.	The failover target did not receive an answer to its ping of the source machine. Eventually, a failover will result. Investigate possible errors (down server, network error, etc.).
5110	Failover	Info.	The %1 script "%2" is still running.	No action required.
5200	Failback	Info..	Failback completed for %1.	No action required.
5201	Failback	Info..	IP address %1 was removed from target machine's %2 adapter.	No action required.
5202	Failback	Error	Unable to Failback properly because IP address %1 was missing a corresponding SubNet Mask.	Contact technical support.
5300	Monitoring	Info.	The following IP address was added to target's monitoring list: %1	No action required.
5301	Monitoring	Info.	The following IP address was removed from target's monitoring list: %1	No action required.
5302	Monitoring	Info.	Drive share information for %1 has been updated on the target machine.	No action required.
5303	Monitoring	Info.	The application monitor script has started successfully.	No action required.
5304	Monitoring	Info	The application monitor script has finished successfully.	No action required.
5305	Monitoring	Warning	The application monitor has found the %1 service stopped.	Double-Take Availability Application Manager will attempt to restart the service.
5306	Monitoring	Warning	The application monitor has restarted the %1 service.	No action required.

ID	Category	Severity	Event Message	Required Response
5307	Monitoring	Error	The application monitor cannot contact the server %1.	Verify the server is running. Verify available network communications with the server.
5400	ARP	Info.	Broadcasted new MAC address %1 for IP address %2.	No action required.
5500	Service	Warning	Could not connect to e-mail server. Check to make sure the SMTP server %1 is available (error code: %2).	Double-Take Backup could not connect to your SMTP server or the username and/or password supplied is incorrect. Verify that SMTP server is available and that you have identified it correctly in your e-mail notification configuration. Also verify that your username and password have been entered correctly.
5501	Service	Warning	E-mail notification could not be enabled (error code: %1).	This alert occurs if there is an unexpected error enabling e-mail notification during 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	Service	Warning	E-mail notification could not be initialized. Check to make sure Internet Explorer 5.0 or later is installed.	E-mail notification requires Internet Explorer 5.0 or later. Verify that you have this version or later installed on the Double-Take Backup server.
5503	Service	Warning	E-mail notification could not be processed. Check to make sure the correct version of SMTPMail.DLL is registered on the system (error code: %1).	If you are using Double-Take Backup 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 Backup server.

ID	Category	Severity	Event Message	Required Response
5504	Service	Warning	Could not load LocalRS.dll (for e-mail notification).	This alert occurs if there is an error loading the resource DLL for the 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	Service	Warning	E-mail could not be sent. Check e-mail settings (error code: %1).	Verify that the e-mail server that you have identified in your e-mail notification configuration is correct.
5506	Service	Warning	One or more required e-mail settings have not been specified (error code: %1).	At a minimum, you must specify the e-mail server, the From and To addresses, and at least one type of event to include.
5507	Service	Warning	E-mail notification could not be initialized. Check to make sure WMI is installed and available (error code: %1).	If you are using Double-Take Backup 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 Backup server.
5508	Service	Warning	An error occurred connecting to the WMI namespace. Check to make sure the Windows Management Instrumentation service is not disabled (error code %1).	This alert occurs if there is an error with the Windows Management Instrumentation (WMI) service. Verify that you have it installed on the Double-Take Backup server and that it is enabled.
5600	Service	Warning	Part or all of the e-mail setting %1 is not in a valid format.	Verify that the include categories and exclude ID list are identified and formatted correctly.

ID	Category	Severity	Event Message	Required Response
7106	RepDrv	Error	The driver was unable to get valid name information from the Filter Manager for the file %2. (Filename may be truncated.) It cannot be replicated. Please contact technical support.	Contact technical support.
7107	RepDrv	Error	The driver was unable to get valid name information from the Filter Manager for a file. It cannot be replicated. Please contact technical support.	Contact technical support.
8100	RepDac	Error	The driver encountered an unrecoverable internal error. Contact technical support. The last Word in the Data window is the internal error code.	Contact technical support.
8192	Resources	Error	Driver failed to allocate Kernel memory. Replication is stopped and server must be rebooted for replication to continue. The last word in the data window is the tag of the allocation that failed.	Reboot the server and contact technical support if this event occurs again.
8192	RepDrv	Error	Kernel memory is exhausted. Replication is stopped. This may have been caused by low system resources.	Reboot the server and contact technical support if this event occurs again.
8193	System	Error	The driver failed to create a thread required for normal operation. This may have been caused by low system resources. Reboot your server and contact technical support if this error occurs again. The last Word in the Data window is the NT status code.	Reboot the server and contact technical support if this event occurs again.
8196	RepDrv	Warning	The maximum amount of memory for replication queuing has been reached. Replication is stopped and memory is being freed.	Contact technical support if this event occurs again.

ID	Category	Severity	Event Message	Required Response
8198	Resources	Warning	The driver registry path could not be saved. The default registry path will be used.	No action required.
8200	Resources	Warning	The driver failed to allocate a buffer for a file name longer than 260 characters. The file will be skipped. The last Word in the Data window is the NT status code.	Reboot the server and contact technical support if this event occurs again.
9000	RepKap	Warning	The driver has failed to process a rename operation. The driver will resend the rename operation. This message is only a warning. If you receive this message repeatedly, contact technical support. The last Word in the Data window is the NT status code.	Contact technical support if this event occurs again.
9100	RepKap	Error	The driver encountered an error opening a file from the service. Check the Event Viewer Application log for additional service information or contact technical support. The last Word in the Data window is the exception code.	Check for related service messages. Contact technical support if this event occurs again.
9101	RepKap	Error	The driver encountered an error reading from the service input buffer. Check the Event Viewer Application log for additional service information or contact technical support. The last Word in the Data window is the exception code.	Check for related service messages. Contact technical support if this event occurs again.
9102	RepKap	Error	The driver encountered an error writing to the service output buffer. Check the Event Viewer Application log for additional service information or contact technical support. The last Word in the Data window is the exception code.	Check for related service messages. Contact technical support if this event occurs again.

ID	Category	Severity	Event Message	Required Response
9103	RepKap	Error	The driver encountered an error writing to the service input buffer. Check the Event Viewer Application log for additional service information or contact technical support. The last Word in the Data window is the exception code.	Check for related service messages. Contact technical support if this event occurs again.
9104	RepKap	Error	The driver encountered an error querying for file security from the service input buffer. Check the Event Viewer Application log for additional service information or contact technical support. The last Word in the Data window is the exception code.	Check for related service messages. Contact technical support if this event occurs again.
9105	RepKap	Error	The driver encountered an error querying for file security from the service input buffer. Check the Event Viewer Application log for additional service information or contact technical support. The last Word in the Data window is the exception code.	Check for related service messages. Contact technical support if this event occurs again.
9106	RepKap	Error	The driver encountered an error writing file security data to the service input buffer. Check the Event Viewer Application log for additional service information or contact technical support. The last Word in the Data window is the exception code.	Check for related service messages. Contact technical support if this event occurs again.
9107	RepKap	Error	The driver encountered an error querying for an allocated range from the service input buffer. Check the Event Viewer Application log for additional service information or contact technical support. The last Word in the Data window is the exception code.	Check for related service messages. Contact technical support if this event occurs again.

ID	Category	Severity	Event Message	Required Response
9108	RepKap	Error	The driver encountered an error querying for an allocated range from the service output buffer. Check the Event Viewer Application log for additional service information or contact technical support. The last Word in the Data window is the exception code.	Check for related service messages. Contact technical support if this event occurs again.
9109	RepKap	Error	The driver encountered an error writing an allocated range to the service input buffer. Check the Event Viewer Application log for additional service information or contact technical support. The last Word in the Data window is the exception code.	Check for related service messages. Contact technical support if this event occurs again.
9110	RepKap	Error	The driver encountered an error querying for a directory from the service input buffer. Check the Event Viewer Application log for additional service information or contact technical support. The last Word in the Data window is the exception code.	Check for related service messages. Contact technical support if this event occurs again.
9111	RepKap	Error	The driver encountered an error querying for a directory from the service output buffer. Check the Event Viewer Application log for additional service information or contact technical support. The last Word in the Data window is the exception code.	Check for related service messages. Contact technical support if this event occurs again.
9112	RepKap	Error	The driver encountered an error writing a directory query to the service input buffer. Check the Event Viewer Application log for additional service information or contact technical support. The last Word in the Data window is the exception code.	Check for related service messages. Contact technical support if this event occurs again.

ID	Category	Severity	Event Message	Required Response
9113	RepKap	Error	The driver encountered an error querying a stream from the service input buffer. Check the Event Viewer Application log for additional service information or contact technical support. The last Word in the Data window is the exception code.	Check for related service messages. Contact technical support if this event occurs again.
9114	RepKap	Error	The driver encountered an error writing a stream query to the service output buffer. Check the Event Viewer Application log for additional service information or contact technical support. The last Word in the Data window is the exception code.	Check for related service messages. Contact technical support if this event occurs again.
9115	RepKap	Error	The driver encountered an error writing a stream query to the service output buffer. Check the Event Viewer Application log for additional service information or contact technical support. The last Word in the Data window is the exception code.	Check for related service messages. Contact technical support if this event occurs again.
9116	RepKap	Error	The driver has failed to close a file handle. If you receive this message repeatedly, contact technical support. The last Word in the Data window is the NT status code.	Contact technical support.
9117	RepKap	Error	The driver encountered an error querying for extended attributes from the service input buffer. Check the Event Viewer Application log for additional service information or contact technical support. The last Word in the Data window is the exception code.	Check for related service messages. Contact technical support if this event occurs again.

ID	Category	Severity	Event Message	Required Response
9118	RepKap	Error	The driver encountered an error writing extended attributes to the service output buffer. Check the Event Viewer Application log for additional service information or contact technical support. The last Word in the Data window is the exception code.	Check for related service messages. Contact technical support if this event occurs again.
9119	RepKap	Error	The driver encountered an error writing extended attributes status to the service input buffer. Check the Event Viewer Application log for additional service information or contact technical support. The last Word in the Data window is the exception code.	Check for related service messages. Contact technical support if this event occurs again.
9120	RepKap	Error	The driver encountered an error querying for file information from the service input buffer. Check the Event Viewer Application log for additional service information or contact technical support. The last Word in the Data window is the exception code.	Check for related service messages. Contact technical support if this event occurs again.
9121	RepKap	Error	The driver encountered an error writing file information to the service output buffer. Check the Event Viewer Application log for additional service information or contact technical support. The last Word in the Data window is the exception code.	Check for related service messages. Contact technical support if this event occurs again.

ID	Category	Severity	Event Message	Required Response
9122	RepKap	Error	The driver encountered an error writing file information status to the service input buffer. Check the Event Viewer Application log for additional service information or contact technical support. The last Word in the Data window is the exception code.	Check for related service messages. Contact technical support if this event occurs again.
9123	RepKap	Error	The driver encountered an error querying for fsctl information from the service input buffer. Check the Event Viewer Application log for additional service information or contact technical support. The last Word in the Data window is the exception code.	Check for related service messages. Contact technical support if this event occurs again.
9124	RepKap	Error	The driver encountered an error writing fsctl information to the service output buffer. Check the Event Viewer Application log for additional service information or contact technical support. The last Word in the Data window is the exception code.	Check for related service messages. Contact technical support if this event occurs again.
9125	RepKap	Error	The driver encountered an error writing fsctl status to the service input buffer. Check the Event Viewer Application log for additional service information or contact technical support. The last Word in the Data window is the exception code.	Check for related service messages. Contact technical support if this event occurs again.
10000	RepHsm	Warning	This message is only a placeholder warning. The last Word in the Data window is the NT status code.	No action required.

ID	Category	Severity	Event Message	Required Response
10000	GeoCluster	Error	Connect failed to node %1 for resource %2. Adding node to reconnect list.	Ensure that GeoCluster is running on all possible owners and that it can communicate on the network selected for mirroring and replication traffic. GeoCluster will try to reestablish a connection using the check unresponsive node interval specified for the resource.
10001	GeoCluster	Info.	Reconnect succeeded to node %1 for resource %2. Will be added as a possible owner when mirror is complete.	No action required.
10002	GeoCluster	Error	Disk check failed on node %1 for resource %2. Removing as a possible owner.	Ensure that GeoCluster is running on all possible owners and that it can communicate on the public network. Also ensure that the disk specified for the resource is functioning correctly on all possible owners.
10003	GeoCluster	Error	Owner %1 of the quorum resource %2 couldn't access the arbitration path %3. Network may be down.	Ensure that the network used to access the arbitration path is up and that the server is operational. Also ensure that the arbitration share path does exist and that the account running the cluster service has write privileges to the share path.
10004	GeoCluster	Warning	Failover of the group %1 is being delayed. Group will be brought online when the target queue is below the limit or the timeout has expired.	No action required.
10005	GeoCluster	Info.	Node %1 is taking ownership of the group %2. The group will be brought online on this node.	No action required.

ID	Category	Severity	Event Message	Required Response
10006	GeoCluster	Warning	The cluster notification thread failed to start on node %1 for resource %2. The resource should be taken offline and brought back online.	Take the resource offline and bring it back online.
10007	GeoCluster	Warning	The user %1 has reverted a snapshot for the %2 resource on node %3.	No action required. The snapshot you selected will be reverted.
10008	GeoCluster	Warning	The user %1 has discarded queued data for the %2 resource on node %3.	No action required. The queue you selected will be discarded.
10009	GeoCluster	Warning	The user %1 is verifying data for the %2 resource on node %3.	A snapshot of the current data has been taken. After you have verified the data, accept or reject the data.
10010	GeoCluster	Warning	The user %1 has rejected the data for the %2 resource on node %3.	No action required. Since the data was rejected, the data has been reverted to the snapshot taken when the data was selected for verification.
10011	GeoCluster	Warning	The user %1 has accepted the data for the %2 resource on node %3.	No action required. The current data will be used.
10012	GeoCluster	Warning	The GeoCluster Replicated Disk resource %1 has been set to validate its data. No data replication is occurring to the remaining nodes in the cluster. Please Accept or Reject the data by right-clicking on the resource and selecting the appropriate option.	Replication has been stopped because of the validation request. Accept or reject the data on the node by right-clicking on the resource and selecting the appropriate option.
10100	RepKap	Error	The driver could not recall a file because it did not have a token for impersonation. The security provider service should set this token. The last Word in the Data window is the exception code.	Contact technical support if this event occurs again.

ID	Category	Severity	Event Message	Required Response
10101	RepKap	Error	The driver could not access the file in the archive bin, due to a failed impersonation attempt. The last Word in the Data window is the exception code.	Contact technical support if this event occurs again.
10102	RepKap	Error	The driver could not recall the file. The last Word in the Data window is the exception code.	Contact technical support if this event occurs again.
11000	Service	Info.	Service has started an archive to %1 (%2) for Replication Set %3, ID: %4	No action required.
11001	Service	Info.	Service has completed an archive to %1 (%2) for Replication Set %3, ID: %4, %5	No action required.
11002	Service	Info.	Service has started a recall from %1 (%2) for Replication Set %3, ID: %4	No action required.
11003	Service	Info.	Service has completed a recall from %1 (%2) for Replication Set %3, ID: %4, %5	No action required.
11004	Service	Warning	Service has failed connection to the RepHSM driver. %1	Reboot the server or manually restart the RepHSM.sys driver.
11005	Service	Warning	Service has aborted the archive operation.	Verify the activation code on the source and target is valid for archiving. Reboot an unlicensed server.
11006	Service	Warning	Service has aborted the archive recall operation.	Verify the activation code on the source and target is valid for archiving. Reboot an unlicensed server.
11007	Service	Warning	Verification has finished with errors. %1	Review the verification log to correct or accept the errors.
11008	Service	Warning	Archive feature is not supported on volume %1.	The source and target must be NTFS for archiving functionality.
11009	Service	Info.	Service has started an archive preview to %1 (%2) for Replication Set %3, ID: %4	No action required.

ID	Category	Severity	Event Message	Required Response
11010	Service	Info.	Service has completed an archive preview to %1 (%2) for Replication Set %3, ID: %4	No action required.
11011	Service	Warning	Service has aborted the archive preview operation.	Verify the activation code on the source and target is valid for archiving. Reboot an unlicensed server.
12000	DTRecall	Info.	The service has started.	This message refers to the Double-Take Recall service. No action required.
12001	DTRecall	Error	The service failed to start.	Check the user name and password for the Double-Take Recall service to ensure validity. Reinstall the software if this event occurs again.
12002	DTRecall	Info.	The service has stopped.	This message indicates a system shutdown or the user stopped the Double-Take Recall service. No action is required.
12003	DTRecall	Error	The service failed to create a stop control event. (Error %1)	Restart the Double-Take Recall service. Reinstall the software if this event occurs again.
12004	DTRecall	Error	RegisterServiceCtrlHandler failed. (Error %1)	Restart the Double-Take Recall service. Reinstall the software if this event occurs again.
12005	DTRecall	Error	Service encountered SetServiceStatus error (Error %1)	Restart the Double-Take Recall service. Reinstall the software if this event occurs again.
12006	DTRecall	Error	Service could not get handle to driver for security update. (Error %1)	The Double-Take Recall service could not connect to the Double-Take Recall archiving driver. Reboot the server and reinstall the software if this event occurs again.

ID	Category	Severity	Event Message	Required Response
12007	DTRecall	Warning	Service failed a periodic security update. (Error %1)	This message refers to the Double-Take Recall service. The operation will be performed every five minutes. Reinstall the software if this event occurs after five minutes.
12288	RepDrv	Error	The driver encountered an error accessing a buffer from the service. Contact technical support. The last Word in the Data window is the exception code.	Contact technical support.
16384	RepDrv	Error	The driver encountered an unrecoverable error. Contact technical support.	Contact technical support
16385	RepDrv	Error	The driver encountered an unexpected internal result. Contact technical support. The last Word in the Data window is the NT status code.	Contact technical support.
16393	RepDrv	Error	The driver encountered an internal error. Contact technical support. The last Word in the Data window is the internal error code.	Contact technical support.
16395	Resources	Error	The driver detected a memory error which may have been caused by a bad driver or faulty hardware. Contact technical support. The last Word in the Data window is the internal error code.	Contact technical support.
16396	Resources	Error	The driver failed to create work queues for normal operation. This may have been caused by low system resources. Reboot the server and contact technical support if this error occurs again. The last Word in the Data window is the NT status code.	Reboot the server and contact technical support if this event occurs again.

ID	Category	Severity	Event Message	Required Response
16400	RepDrv	Info.	RepDrv has encountered an unexpected condition, usually caused by low kernel memory. Unless otherwise mentioned, this event has already been handled and your data remains protected. If you continue to receive these events or have further questions please contact tech support.	No action required.

E-mailing source event messages

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

1. To configure e-mail notification for source event message, you will need to use the Replication Console. Select **Start, Programs, Double-Take, Double-Take Replication Console**.
2. Right-click the Double-Take Backup source server on the left pane of the Replication Console.
3. Select **Properties**.
4. Select the **E-mail Notification** tab.
4. Select **Enable notification**.

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

5. Specify your e-mail settings.
 - **Mail Server (SMTP)**—Specify the name of your SMTP mail server.
 - **Log on to SMTP Server**—If your SMTP server requires authentication, enable **Log on to SMTP Server** and specify the **Username** and **Password** to be used for authentication. Your SMTP server must support the LOGIN authentication method to use this feature. If your server supports a different authentication method or does not support authentication, you may need to add the Double-Take Backup 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 Backup e-mail message. The address is limited to 256 characters.
 - **Send To**—Specify the e-mail address that each Double-Take Backup 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 Backup e-mail message. This will help distinguish Double-Take Backup 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**. (By default, the test will be run from the machine where the Replication Console is running.) If desired, you can send the test message to a different e-mail address by selecting **Send To** and entering a comma or semicolon separated list of addresses. Modify the message text up to 1024 characters, if necessary. Click **Send** to test the e-mail notification. The results will be displayed in a message box. Click **OK** to close the message and click **Close** to return to the **E-mail Notification** tab.

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

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

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

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

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

Statistics

Statistics logging is the process of taking snapshots of Double-Take Backup 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

```
=====
0/11/09 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 Backup 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	<ul style="list-style-type: none"> • -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	<ul style="list-style-type: none"> • 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	<ul style="list-style-type: none"> • 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 Backup 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.

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

- The statistic labels will be slightly different in the ASCII file than in the following table.
- The statistics will appear in a different order in the ASCII file than in the following table.

- The statistics in the Target Category in the following table are not included in the ASCII file.
- The Kernel statistic Target Open Handles is not included in the ASCII file.
- The ASCII file contains a Managed Pagefile Alloc statistic which is no longer used.

Category	Statistic	Description
Date/Time Stamp		The date and time that the snapshot was taken. This is the date and time that each statistic was logged. By default, these are generated once a second, as long as there are statistics being generated. If mirroring/replication is idle, then DTStat will be idle as well.
SystemAllocator	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	FailedLogins	The number of failed login attempts
Kernel	SourceState	<ul style="list-style-type: none"> • 0—Source is not running • 1—Source is running without the replication driver • 2—Source is running with the replication driver
Kernel	TargetState	<ul style="list-style-type: none"> • 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 Backup 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).

Category	Statistic	Description
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

Category	Statistic	Description
Target	OpsDropped	The number of operations skipped during a difference mirror. During a difference mirror, if Double-Take Backup 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 Backup. 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 Backup. 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.

Category	Statistic	Description
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	<p>The state of the active connection</p> <ul style="list-style-type: none"> • 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). <p>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).</p>
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

Category	Statistic	Description
Connection	conOpsTx	The number of operations transmitted to the target. This is the total number of operations that Double-Take Backup 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 Backup 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 Backup 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	conMirrorBytesRemain	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

Category	Statistic	Description
Connection	conTaskCmdsQueued	The number of task commands that have been queued on the source
Connection	conTasksSucceeded	The number of task commands that have succeeded on the source
Connection	conTasksFailed	The number of task commands that have failed on the source
Connection	conTasksIgnored	The number of task commands that have been ignored on the source

Performance Monitor

Performance Monitor is the Windows graphical tool for measuring performance. It provides charting, alerting, and reporting capabilities that reflect both current activity and ongoing logging. Double-Take Backup statistics are available through the Performance Monitor.

- [Monitoring Performance Monitor statistics](#)
- [Performance Monitor statistics](#)

Monitoring Performance Monitor statistics

1. To access the Performance Monitor, select **Start, Programs, Administrative Tools, Performance**.
2. Specify the data to monitor by right-clicking and selecting **Add** or using the **Add** button on the toolbar.
3. Choose one of the following Double-Take Backup Performance Objects.
 - Double-Take Connection
 - Double-Take Kernel
 - Double-Take Security
 - Double-Take Source
 - Double-Take Target
4. Select the statistics you want to monitor, and click **Add**.

For additional information and details on the Performance Monitor, see your Windows reference guide.

Performance Monitor statistics

The following table identifies the Double-Take Backup Performance Monitor statistics.

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

Object	Statistic	Description
Connection	Bandwidth Limit	The amount of bandwidth that may be used to transfer data
Connection	Bytes in disk queue	The number of bytes in the source disk queue
Connection	Bytes in replication queue	The number of replication bytes in the source queue
Connection	Bytes in mirror queue	The number of mirror bytes in the source queue
Connection	Bytes received	The number of bytes received by the target since the last Performance Monitor refresh
Connection	Bytes transferred	The number of bytes transmitted from the source
Connection	Compressed bytes transferred	The number of compressed bytes transmitted from the source
Connection	Operations in acknowledgement queue	The number of operations waiting in the source acknowledgement queue
Connection	Operations in command queue	The number of operations waiting in the source command queue
Connection	Operations in mirror queue	The number of mirror operations in the source queue
Connection	Operations in replication queue	The number of replication operations in the source queue
Connection	Operations received	The number of operations received by the target since the last Performance Monitor refresh
Connection	Operations resent	The number of operations re-sent since the last time the Double-Take service was restarted on the source
Connection	Operations transmitted	The number of operations transmitted from the source
Connection	Task commands queued	The number of task commands queued on the source
Connection	Task commands submitted	The number of task commands submitted on the source
Connection	Tasks failed	The number of task commands that have failed to execute on the source

Object	Statistic	Description
Connection	Tasks ignored	The number of task commands that have been ignored on the source
Connection	Tasks succeeded	The number of task commands that have succeeded on the source
Kernel	Activation code failures	The number of activation code failures when loading the source or target, since the last time the Double-Take service was restarted on the source
Kernel	Double-Take queue memory usage	The amount of system memory in use by the Double-Take Backup queue
Kernel	Driver Queue Percent	The amount of throttling calculated as a percentage of the stop replicating limit
Kernel	Failed mirror operations	The number of mirror operations on the source that failed due to an error reading the file from the disk
Kernel	Failed replication operations	The number of replication operations on the source that failed due to an error reading the file from the disk
Kernel	Mirror Kbytes generated	The number of mirror kilobytes transmitted from the source
Kernel	Mirror operations generated	The number of mirror operations transmitted from the source
Kernel	Open Target Handles	The number of handles currently open on the target.
Kernel	Replication Kbytes generated	The number of replication kilobytes generated on the source by the file system driver
Kernel	Replication operations generated	The number of replication operations generated on the source by the file system driver
Security	Failed logins	Number of failed login attempts since the last time the Double-Take service was restarted
Security	Successful logins	Number of successful login attempts since the last time the Double-Take service was restarted
Source	Auto disconnects	The number of automatic disconnects since the last time the Double-Take service was restarted on the source
Source	Auto reconnects	The number of automatic reconnects since the last time the Double-Take service was restarted on the source

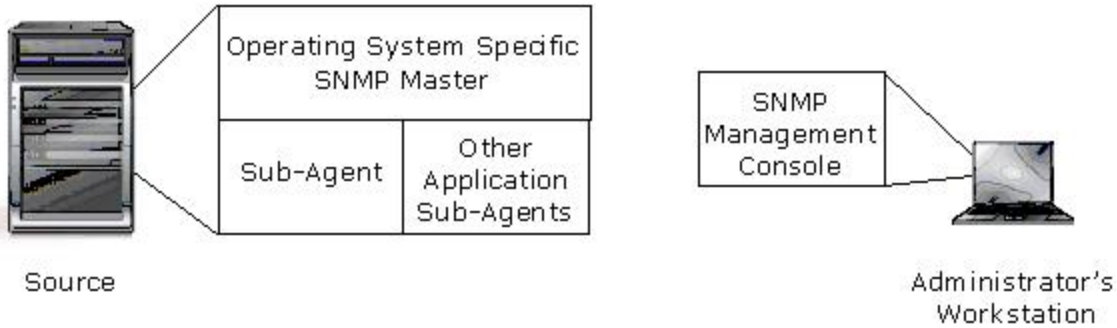
Object	Statistic	Description
Target	Bytes in Disk Queue	The number of bytes in the target disk queue
Target	Bytes in Queue	The number of bytes in the system memory and disk queues
Target	Mirror operations received	The number of mirror operations received on the target
Target	Operations received	The number of operations received on the target
Target	Ops Dropped	The number of operations dropped on the target since the last time the Double-Take service was restarted on the target
Target	Ops Remaining	The number of operations on the target remaining to be applied
Target	Orphan Bytes	The number of orphan bytes removed from the target
Target	Orphan Directories	The number of orphan directories removed from the target
Target	Orphan Files	The number of orphan files removed from the target
Target	Retries	The number of retries performed on the target since the last time the Double-Take service was restarted on the target
Target	Tasks failed	The number of task commands that have failed on the target.
Target	Tasks ignored	The number of task commands that have been ignored on the target
Target	Tasks succeeded	The number of task commands that have succeeded on the target

SNMP

SNMP, Simple Network Management Protocol, is the Internet's standard for remote monitoring and management of hosts, routers and other nodes and devices on a network. Double-Take Backup provides an SNMP sub-agent that can be managed from an SNMP Management Console.

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

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



Configuring SNMP on your server

SNMP must be installed on a server before Double-Take Backup in order for the Double-Take Backup SNMP components to be added during the Double-Take Backup installation. If SNMP is installed on a server after Double-Take Backup is installed, run a repair install to install the SNMP components.

The Double-Take Backup .mib file will need to be loaded into your SNMP Management Console. Depending on the type of console you are using, this process might include compiling the .mib file. Reference your SNMP Management Console documentation for additional information.

SNMP traps

The following table lists the Double-Take Backup SNMP traps.

Object Type	Trap	Description
Kernel	dttrapKernelStarted	Double-Take Backup has started
Kernel	dttrapKernelStopped	Double-Take Backup has stopped
License	dttrapLicenseViolationStartingSource	The source cannot be started due to a license violation
License	dttrapLicenseViolationOnNetwork	A Double-Take Backup serial number conflict was identified on the network
Source	dttrapSourceStarted	Double-Take Backup source component has started
Source	dttrapSourceStopped	Double-Take Backup source component has stopped
Target	dttrapTargetStarted	Double-Take Backup target component has started
Target	dttrapTargetStopped	Double-Take Backup target component has stopped

Object Type	Trap	Description
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 Backup 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 Backup to shutdown
Replication	dttrapReplicationStart	Replication has started
Replication	dttrapReplicationStop	Replication has stopped
Mirroring	dttrapMirrorStart	Mirroring has started
Mirroring	dttrapMirrorStop	Mirroring has stopped

Object Type	Trap	Description
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 Backup SNMP statistics.

Object Type	Statistic	Description
General	dtUpTime	Time in seconds since Double-Take Backup was last started
General	dtCurrentMemoryUsage	Amount of memory allocated from the Double-Take Backup 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

Object Type	Statistic	Description
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	<ul style="list-style-type: none"> • 0—Source is not running • 1—Source is running without the replication driver • 2—Source is running with the replication driver.
Target	dtTargetState	<ul style="list-style-type: none"> • 0—Target is not running • 1—Target is running
Target	dtRetryCount	The number of file operations that have been retried
Target	dtOpsDroppedCount	The number of file operations that have failed and will not be retried
Security	dtLoginCount	The number of successful logins
Security	dtFailedLoginCount	The number of unsuccessful logins
Connection	dtConnectionCount	The number of active connections between machines
Connection	dtconIpAddress	The IP address of the connected machine. If at the source, then the IP address of the target. If at the target, then the IP address of the source.
Connection	dtconConnectTime	The duration of time since the connection was first established

Object Type	Statistic	Description
Connection	dtconState	<p>The state of the active connection</p> <ul style="list-style-type: none"> • 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). <p>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).</p>
Connection	dtconOpsInCmdQueue	The number of operations (create, modify, or delete) in the retransmit queue on the source
Connection	dtconOpsInAckQueue	The number of operations (create, modify, or delete) waiting for verification acknowledgements from the target
Connection	dtconOpsInRepQueue	The number of replication operations (create, modify, or delete) in the queue
Connection	dtconOpsInMirQueue	The number of mirror operations (create, modify, or delete) in the queue
Connection	dtconBytesInRepQueue	The number of bytes in the replication queue
Connection	dtconBytesInMirQueue	The number of bytes in the mirror queue
Connection	dtconOpsTx	The total number of operations (create, modify, or delete) transmitted to the target
Connection	dtconBytesTx	The total number of bytes transmitted to the target
Connection	dtconBytesCompressedTx	The total number of compressed bytes transmitted to the target
Connection	dtconOpsRx	The total number of operations (create, modify, or delete) received from the target
Connection	dtconBytesRx	The total number of bytes received from the target
Connection	dtconResentOpCount	The number of operations that were resent because of acknowledgement errors

Error codes

The following table contains error codes that you may see in the various user interfaces or in log files.

Error Code	Description
-1	Unknown error code (generated when a command failed but the failure is not linked to a pre-defined error code)
-101	Invalid parameter was supplied
-102	Command is not a valid or the syntax is incorrect
-103	Double-Take Backup source module is not loaded
-104	No Double-Take Backup source identified
-105	Double-Take Backup target module is not loaded
-106	Connection already established
-107	Connection does not exist
-108	Mirror currently active
-109	Server does not exist or could not be located
-110	Server is not responding
-111	Double-Take Backup is running
-112	Unknown connection error
-113	Mirror already active
-114	Date is invalid - valid format is mm/dd/yy
-115	Time is invalid - valid format is hh:mm
-116	Invalid option supplied
-117	Mirror is not paused
-118	Connection is not paused
-119	Connection does not exist
-120	Connection already connected
-121	Mirror is not running
-122	Replication set exists

Error Code	Description
-123	Replication set does not exist
-124	No replication set has been selected
-125	Connection is replicating
-126	Connection is not replicating
-127	Replication set is enabled
-128	Schedule is not defined
-129	Replication set is changed
-130	Replication set is in use
-131	No Double-Take Backup target identified
-132	Memory is low
-133	Memory is sufficient
-134	Replication is pending
-135	Invalid option supplied
-136	Replication set rule does not exist
-137	Mirror queue is full
-138	Insufficient security access
-139	Schedule command is invalid
-140	Source path is invalid
-141	Replication set is not changed
-142	Insufficient source security access
-143	Invalid statistics file
-144	Replication set not saved
-145	Connection failed
-146	Cleaner option is not enabled
-147	Target mirror capacity high threshold is met
-148	Target mirror capacity low threshold is met
-149	New option applied

Error Code	Description
-150	Target is restarted
-151	Replication is out of memory
-152	Write access is blocked on the volume
-153	This error code could be one of two errors. 1) Compression level is not supported, or server does not support compression 2) Transmission is paused
-154	Transmission is active
-155	Target does not support the command
-156	Command conversion to accommodate a different Double-Take Backup version has failed
-157	Incompatible source and target Double-Take Backup versions
-158	Incompatible source and target operating system versions
-159	NAS server to non-NAS server is not a supported configuration
-160	Target module is not loaded
-161	Operation or command is not supported
-162	Target is paused
-163	Target is pending
-164	Target is active
-165	Target is retrying operations
-166	Target is no longer retrying operations
-167	Restore required state is unknown
-168	Not a valid failover source
-169	Failover login failed
-170	Feature is not supported
-171	Command is not supported
-172	Target queue log file error
-173	Target disk is full
-174	Target disk has sufficient disk space
-175	Error reading from or writing to the queue log file

Error Code	Description
-176	Memory-based queue is in use
-177	Disk-based queue is in use
-178	Restore is required
-179	ID the driver supplied to the service is invalid
-180	Child path is blocked
-181	Parent path is blocked
-182	Target path blocking is disabled
-183	Connection ID specified is invalid
-184	No command objects are in the queue
-185	Target is discarding operations from the target queue
-186	Target is not discarding operations from the target queue
-187	Schedule is paused
-188	Schedule is resumed
-189	Target state has changed
-190	Target name has changed
-201	Monitor name exists
-202	Monitor name does not exist
-203	Monitor configuration exists
-204	Monitor configuration does not exist
-205	Monitor configuration is in use
-206	Monitor configuration is not in use
-207	Source is online
-208	Source is offline
-209	Server is not failed over
-210	Server is failed over
-211	Server is not being monitored
-212	Failback is in progress

Error Code	Description
-213	IP address placeholders on the target are unavailable
-214	Target NIC was not found
-215	Source module is not loaded
-216	Failed to set the source state
-217	Unable to ping source
-218	Invalid argument
-219	Recovery is busy
-220	Invalid command
-221	Recovery is started
-222	Script failed to start
-223	Script timeout met
-224	No replication timeout met - connection is bad
-225	Invalid path
-226	Kernel module is not loaded
-2201	Error communicating with e-mail server
-2202	Error connecting to e-mail server
-2203	E-mail notification is disabled
-2204	E-mail notification is enabled
-2205	E-mail notification requires Internet Explorer version 5.0 and WMI
-2206	E-mail notification requires Internet Explorer version 5.0
-2207	Error sending e-mail
-2208	Error sending test e-mail
-2209	WMI error connecting to e-mail server
-2210	E-mail notification requires WMI
-2211	Event Viewer settings for e-mail notification are invalid
-2212	E-mail notification setting is invalid
-2213	E-mail notification address exists

Error Code	Description
-2214	E-mail notification alert ID is invalid
-2215	E-mail notification format is invalid
-2216	E-mail notification address does not exist
-2217	E-mail notification address notification list is empty
-2218	E-mail warning is not set
-2219	E-mail test warning is not set
2200	E-mail notification is functioning properly
-2301	Bandwidth limiting time exists
-2302	Bandwidth limiting name exists
-2303	Bandwidth limit not found
-2304	Bandwidth limit day is invalid
-2305	Bandwidth limit lable is invalid
-2401	Snapshot module is not loaded
-2402	Error reading the snapshot .dll
-2403	Snapshot not found
-2404	No snapshot connections found
-2501	Full-server functionality is disabled
-2502	No full-server interface available

Cargo archiving

Before you begin archiving, you must [configure specific archiving security](#). Once that is complete, you may want to generate a preview report to help you plan archiving and fine-tune your archive settings. When you are ready to begin archiving, you can archive files manually or you can schedule archiving for periodic intervals.

- [Running an archive preview report](#)
- [Manually archiving data](#)
- [Archiving data on a schedule](#)
- [Choosing a storage location for archived files on the repository server](#)
- [Recalling archived files](#)

Note: The archiving functionality is currently only available in the Replication Console. In the future, the archiving functionality will be made available in the Double-Take Backup console. By default, the Replication Console is installed with the Double-Take Backup installation. To open the console, select **Start, Programs, Double-Take, Double-Take Replication Console**.

Archiving is not compatible with [TimeData](#) continuous data protection. If you use one, do not use the other.

Configuring archiving security

Before you can use Double-Take Backup archiving, you must establish a specific security configuration. This is a six step process.

1. Confirm Double-Take Backup is installed on both the source and repository server.
2. Create a new service account. (A service account is a user account that is created explicitly to provide security context for a service.) Follow steps a through f to create a new service account.
 - a. From **Active Directory Users and Computers**, create a new user.
 - b. Enter a descriptive name for the first and last name and modify the full name as desired.
 - c. Specify a **User logon name**.
 - d. Specify and confirm a **Password** for the account.
 - e. Specify your password settings. The settings you select may be dependent on your company's security policies. Keep in mind the following caveats for password settings for a service account.
 - **User must change password at next logon**—The Double-Take Software recommendation is to disable this setting. The Double-Take Recall service will not be able to start if this option is enabled because the service will be waiting on the required logon change.
 - **User cannot change password**—The Double-Take Software recommendation is to enable this setting. If this setting is enabled, you will not have to worry about updating the credentials in the Double-Take Recall service. If you cannot select this option because of company security policies and the password on the account is changed, the Double-Take Recall service will no longer have valid credentials. File recalls will not function until the Double-Take Recall service credentials are updated to the new password.
 - **Password never expires**—The Double-Take Software recommendation is to enable this setting. If this setting is enabled, you will not have to worry about updating the credentials in the Double-Take Recall service. If you cannot select this option because of company security policies and the password on the account expires, the Double-Take Recall service will no longer have valid credentials. File recalls will not function until the account is reset and the Double-Take Recall service credentials are updated.

-
- **Account is disabled**—The Double-Take Software recommendation is to disable this setting. Because the word disable is part of the option name, the recommendation can be confusing. You want to enable the account, which means this option should not have a checkmark. If the option does have a checkmark, meaning the account would be disabled, file recalls will not function.
- f. Finish the account creation wizard.
 3. Add the new service account to the local **Administrators** group on the source and repository server.
 4. Add the new service account to the local **Double-Take Recall** group on the source and repository server.
 5. Modify the Double-Take Recall service on the source and repository server to use the new service account.
 - a. From the services applet (**Administrative Tools, Services**) right-click the Double-Take Recall service and select **Properties**.
 - b. On the **Log On** tab, select **This account** and specify the new user and the password.
 - c. On the **General** tab, select **Startup type**, and select **Automatic**.
 - d. Click **OK** to save the changes.
 6. Start the Double-Take Recall service on the source and repository server. If the service is already running, you must stop and restart it so that it uses the user account and group modifications you just made.

If you have an active Double-Take Backup job, you can validate your security configuration using the Replication Console.

1. Open the Replication Console by selecting **Start, Programs, Double-Take, Double-Take Replication Console**.
 2. Double-click your source machine in the left pane of the Replication Console to log on to it.
 3. With the source highlighted in the left pane, locate in the right pane the Double-Take protection job in the **Replication Set** column. The name will be created from the source and repository server names and the keyword **Protection** or **Data Protection**.
- Note
4. Right-click the protection job in the right pane of the Replication Console and select **Connection Manager**.
 5. Select the **Archive Options** tab.
 6. Click **Validate**.

The validation checks to see if the correct service is running and if it has an account in the correct security group. It also confirms if the repository server has shared the archive bin correctly and if network communications are functioning properly. Any success or failure messages will be included in the Double-Take Backup log file on the source.

Running an archive preview report

To help you plan archiving and fine-tune your archive settings, you can generate a preview report. The archiving preview report does not actually archive any data but reports how much data would be archived based on the settings you select.

1. Open the Replication Console by selecting **Start, Programs, Double-Take, Double-Take Replication Console**.
2. Double-click your source machine in the left pane of the Replication Console to log on to it.
3. With the source highlighted in the left pane, locate in the right pane the Double-Take protection job in the **Replication Set** column. The name will be created from the source and repository server names and the keyword **Protection** or **Data Protection**.

4. Right-click the protection job in the right pane of the Replication Console and select **Archiving, Preview**.
5. Select the archiving options that you would like to perform.



- **Archive files inactive for**—This option specifies the age in days, based on the last accessed time, of the files that you would want archived. Any file older than the age specified will be included in the preview report.
 - **Archive files when size meets or exceeds**—This option specifies the size, in KB, of the files that you want archived. Any file equal to or larger than the specified size will be archived.
6. If desired, specify if you want to calculate the size of the replication set (the protection job) prior to starting the archive. If enabled, this will display the archiving percentage complete in the **Mirror Status** column of the Replication Console.
 7. Click **Preview** to generate the report. While the preview report is being generated, you will see the **Mirror Status** column change to **Archiving**. When the preview report is complete, the **Mirror Status** column will change back to **Idle**.
 8. After the preview report has been generated, locate and open the file from the location where Double-Take Backup is installed. The file name uses the name of the protection job followed by ArchivePreview.txt. For example, if your protection job is called alpha to beta data protection, then the preview report would be called alpha to beta data protection ArchivePreview.txt.

Note If you want to change the file name used to generate the preview report, right-click the server from the left pane of the Replication Console and select **Properties**. On the **Logging** tab, modify the **Archive Preview Filename**. On this tab, you can also modify the maximum size of the preview report.

9. Review the preview report. At the top of the report, you will find overview information, including the archive criteria applied to the preview. In the center of the report will be a list of all of the files that would be archived based on the selected criteria. At the bottom of the report, you will find summary statistics.
10. If desired, repeat the preview report using different archiving criteria. The preview report is replaced each time a preview is run, so rename the file to a different file name if you want to keep the report.

Manually archiving data

Archiving can be performed manually anytime after an initial mirror and when another mirror is not in progress. You can archive individual files and folders, or you can archive groups of files based on archive criteria.

Archiving files individually

To archive files or folders individually, use the following steps.

1. Open the Replication Console by selecting **Start, Programs, Double-Take, Double-Take Replication Console**.
2. Double-click your source machine in the left pane of the Replication Console to log on to it.
3. Expand the tree under the source and locate the Double-Take protection job. The name will be created from the source and repository server names and the keyword **Protection** or **Data Protection**.
4. Expand the tree under the job name.
5. Highlight the root of the volume under the job name that contains the files or folders that you want to archive.
6. In the right pane of the Replication Console, drill down through the volume to locate the files or folders that you want to archive.
7. Right-click on the files or folders and select **Archive**. The right-click menu will not be available if the files or folders are not part of your protection job.

Archiving groups of files based on criteria

To archive groups of files based on archive criteria, use the following steps.

1. Open the Replication Console by selecting **Start, Programs, Double-Take, Double-Take Replication Console**.
2. Double-click your source machine in the left pane of the Replication Console to log on to it.
3. With the source highlighted in the left pane, locate in the right pane the Double-Take protection job in the **Replication Set** column. The name will be created from the source and repository server names and the keyword **Protection** or **Data Protection**.
4. Right-click the protection job in the right pane of the Replication Console and select **Archiving, Preview**.
5. Select the archiving options that you would like to perform. Both archiving criteria must be met for a file to be archived.



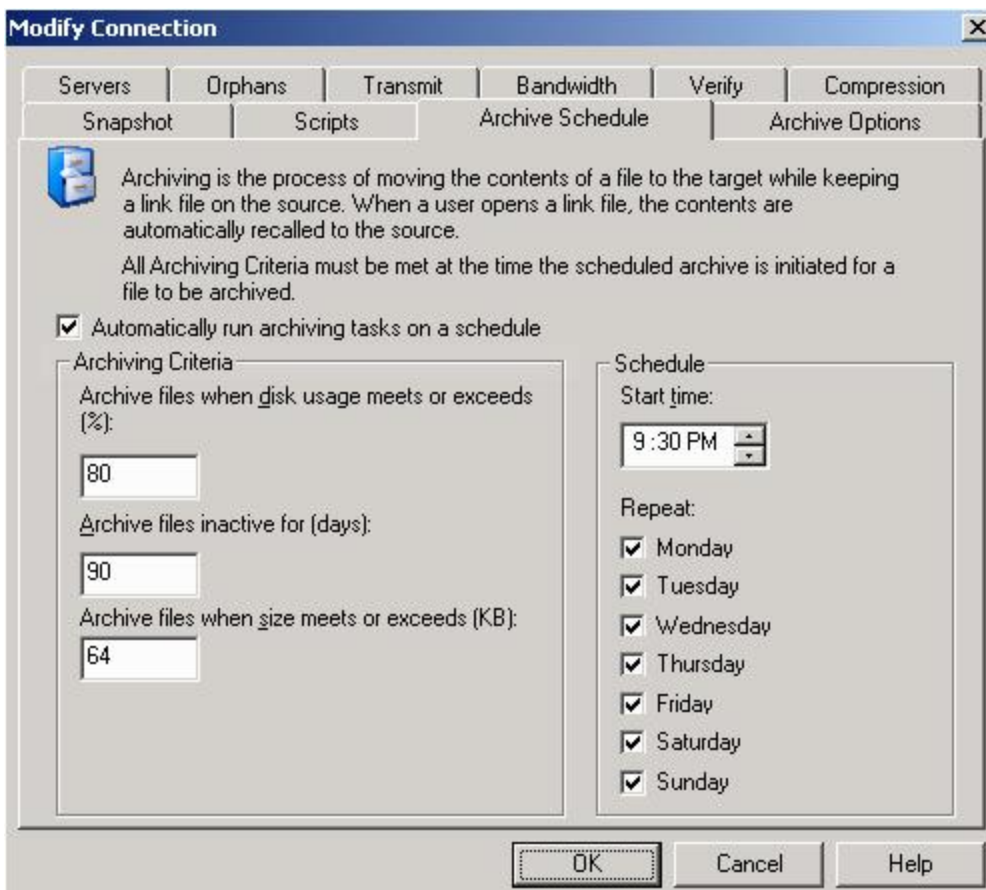
- **Archive files inactive for**—This option specifies the age in days, based on the last accessed time, of the files that you want archived. Any file older than the age specified will be archived.
 - **Archive files when size meets or exceeds**—This option specifies the size, in KB, of the files that you want archived. Any file equal to or larger than the specified size will be archived.
6. If desired, specify if you want to calculate the size of the replication set (the protection job) prior to starting the archive. If enabled, this will display the archiving percentage complete in the **Mirror Status** column in the right pane of the Replication Console.
 7. Click **Archive** to start the process.

While the archiving is running, you will see the **Mirror Status** column in the Replication Console change to **Archiving**. During this time, you can pause or resume archiving by right-clicking the protection job and selecting **Archiving, Pause** or **Archiving, Resume**. If you need to stop the archiving process, select **Archiving, Stop**.

Archiving data on a schedule

Archiving can be scheduled to occur automatically at periodic intervals.

1. Open the Replication Console by selecting **Start, Programs, Double-Take, Double-Take Replication Console**.
2. Double-click your source machine in the left pane of the Replication Console to log on to it.
3. With the source highlighted in the left pane, locate in the right pane the Double-Take protection job in the **Replication Set** column. The name will be created from the source and repository server names and the keyword **Protection** or **Data Protection**.
4. Right-click the protection job in the right pane of the Replication Console and select **Connection Manager**.
5. Select the **Archive Schedule** tab.



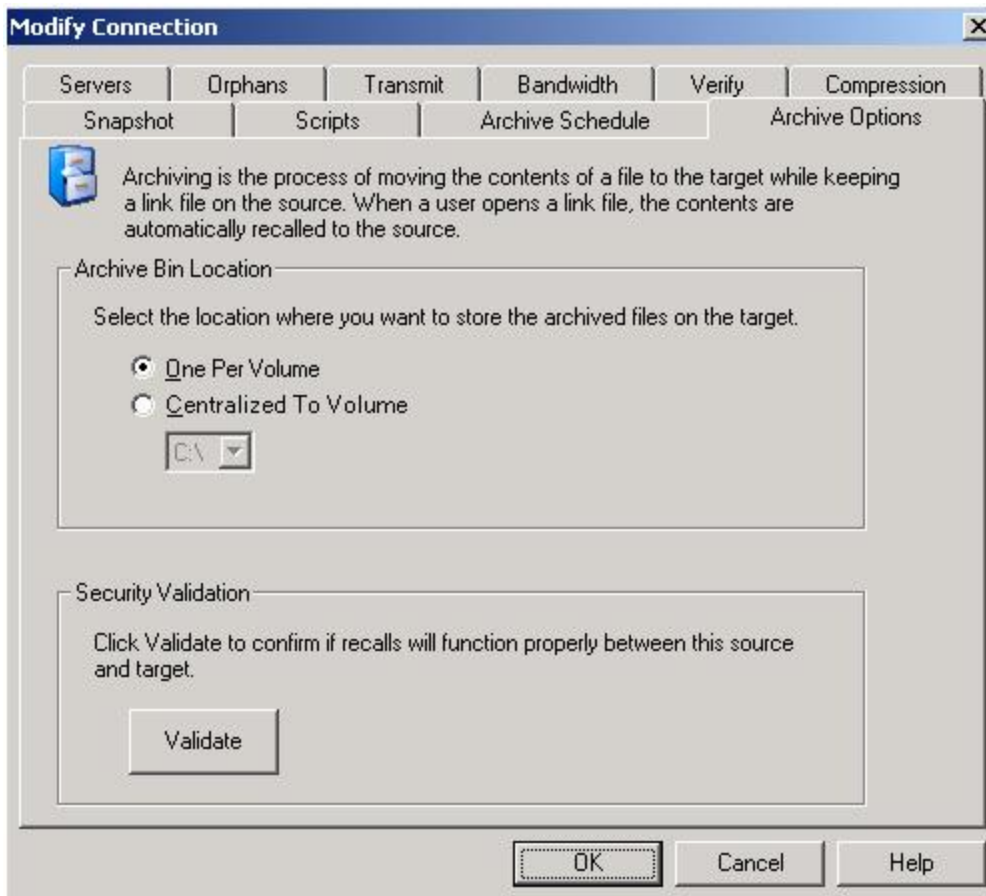
-
6. To enable scheduled archiving, enable **Automatically run archiving tasks on a schedule**.
 7. Specify the criteria that you want to apply to your protection job. All **Archiving Criteria** must be met at the time the scheduled archive is initiated for a file to be archived.
 - **Archive files when disk usage meets or exceeds**—This option is a first level filter. If disk usage exceeds the specified percentage, Double-Take Backup will check the remaining criteria to determine if any files should be archived. If disk usage is not above the specified percentage, no archiving will occur.
 - **Archive files inactive for**—This option is a second level filter. It specifies the age in days, based on the last accessed time, of the files that you want archived. If disk usage has been exceeded, any file older than the age specified will be archived.
 - **Archive files when size meets or exceeds**—This option is also a second level filter. It specifies the size, in KB, of the files that you want archived. Any file equal to or larger than the specified size will be archived.
 6. Specify the archiving schedule that you want to use. Specify the **Start time** and the days of the week to run the archiving process.
 7. Click **OK** to save the settings.

When the archive process is initiated by the schedule, the **Mirror Status** column in the right pane of the Replication Console will change to **Archiving**.

Note: If you disable the scheduled archive after it has already archived files, the archived files will remain as link files on the source until they are recalled. No additional files will be archived unless you enable the schedule again or manually initiate an archive.

Choosing a storage location for archived files on the repository server

1. Open the Replication Console by selecting **Start, Programs, Double-Take, Double-Take Replication Console**.
2. Double-click your source machine in the left pane of the Replication Console to log on to it.
3. With the source highlighted in the left pane, locate in the right pane the Double-Take protection job in the **Replication Set** column. The name will be created from the source and repository server names and the keyword **Protection** or **Data Protection**.
4. Right-click the protection job in the right pane of the Replication Console and select **Connection Manager**.
5. Select the **Archive Options** tab.



6. Select your **Archive Bin Location**. You have the choice of putting your archived files on the repository server on the same volume location as they are on the source, or you can place all of the archived files from all volumes on the source on one volume on the repository server. Specify **One Per Volume** to use the same volume structure on the source and repository server. Specify **Centralized to Volume** and specify a volume to use one volume on the repository server for archived files from all volumes on the source.
7. Click **OK** to save the settings.

Recalling archived files

When a file has been archived, the link file is still accessible on the source. The link file has an offline file attribute associated with it. When an end-user attempts to access the link file, the contents of the file are recalled from the repository server back to the source. The file opens to the end-user as if it had always been fully on the source.

If a link file is moved to a different volume on the source, the file will be recalled and moved. If a link file is moved within the same volume on the source, the file is not recalled, only moved.

There may be times when you want to recall some or all of the files from the repository server back to the source. For example, you should recall all of the files before changing to a new repository server to protect your source or if you want to rearchive your files using different archive settings.

Recalling files individually

To recall files or folders individually, use the following steps.

-
1. Open the Replication Console by selecting **Start, Programs, Double-Take, Double-Take Replication Console**.
 2. Double-click your source machine in the left pane of the Replication Console to log on to it.
 3. Expand the tree under the source and locate the Double-Take protection job. The name will be created from the source and repository server names and the keyword **Protection** or **Data Protection**.
 4. Expand the tree under the job name.
 5. Highlight the root of the volume under the job name that contains the files or folders that you want to archive.
 6. In the right pane of the Replication Console, drill down through the volume to locate the files or folders that you want to recall.
 7. Right-click on the files or folders and select **Recall**. The right-click menu will not be available if the files or folders are not part of your protection job.

Recalling all files

To recall all of the files from the repository server, use the following steps.

1. Open the Replication Console by selecting **Start, Programs, Double-Take, Double-Take Replication Console**.
2. Double-click your source machine in the left pane of the Replication Console to log on to it.
3. With the source highlighted in the left pane, locate in the right pane the Double-Take protection job in the **Replication Set** column. The name will be created from the source and repository server names and the keyword **Protection** or **Data Protection**.
4. Right-click the protection job in the right pane of the Replication Console and select **Archiving, Recall All**.
5. If desired, specify if you want to calculate the size of the replication set (the protection job) prior to starting the recall. If enabled, this will display the recall percentage complete in the **Mirror Status** column in the right pane of the Replication Console.
6. Click **Recall All Files** to begin the recall.

Using firewalls

Double-Take Backup can be used over a wide area network (WAN) through firewalls, but it does not support Network Address Translation (NAT) configurations.

If your servers are on opposite sides of a firewall, you will need to configure your hardware to accommodate Double-Take Backup communications. You must have the hardware already in place and know how to configure the hardware ports. If you do not, see the reference manual for your hardware.

Double-Take Backup ports

Double-Take Backup port information is stored in configuration files. By default, ports 80, 6320, 6330, and 6340 are used. If you are using TimeData continuous data protection, the default ports are 8080, 17001, 17101, 17201, and 17301. If desired you can modify some of these ports. However, not all of the ports can be modified.

If you have to modify which ports are used, you will have to modify two configuration files. Be careful when modifying these configuration files. If you modify them incorrectly, Double-Take Backup may no longer work correctly.

1. Open the file JobMgr.dll.config in your Double-Take Backup installation directory. By default, this is \Program Files\Double-Take Software\Double-Take. You may need to associate the .config file extension with Notepad or another text editor.
2. Locate the following lines near the top of the configuration file.

```
<add key="IdentificationPort" value="6320"/>
<add key="FallbackIdentificationPort" value="1100"/>
<add key="LivewirePort" value="6340"/>
<add key="VRAPort" value="6330"/>
<add key="TimeDataWebServicePort" value="8080"/>
```
3. Change the port numbers as necessary. Do not change any other information.
4. Save the configuration file.
5. Repeat steps 1 through 4 with the file BackupConsole.exe.config, using the same port numbers for each key as you used in the JobMgr.dll.config file.
6. Repeat steps 1 through 5 for each Double-Take Backup server. The ports must match between all Double-Take Backup servers and clients.
7. If you changed the IdentificationPort in the configuration files, you must also change a port in the Replication Console.
 - a. Select **Start, Programs, Double-Take , Double-Take Replication Console**.
 - b. Locate your server in the server tree in the left pane of the Replication Console.

Note: If your server is not listed, select **Insert, Server**. Type the machine name or IP address and click **OK**.

- c. Double-click the server to log on to it.
- d. Right-click the server in the left pane of the Replication Console and select **Properties**.
- e. On the **Network** tab, modify the **Communications Port** to match the IdentificationPort in the configuration files.
- f. Click **OK**.
- g. Repeat steps b through f for all of your Double-Take Backup servers.

All of your Double-Take Backup servers and clients must have the same port configuration.

Microsoft Windows ports

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

Double-Take Backup also relies on other Microsoft Windows ports.

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

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

Hardware ports

You need to configure your hardware so that all of the above ports for both Double-Take Backup and Microsoft Windows are open. Since communication occurs bi-directionally, make sure you configure both incoming and outgoing traffic.

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

Index

A

ACLs (access control list)	10
activation code	21
alternate data streams	10
archiving	See Cargo
Atom	2

B

bandwidth	30, 57
-----------	--------

C

Cargo	
files	126
overview	2, 4, 123
preview report	124
recalling	129
requirements	5, 9
schedule	127
security	123
storage location	128
compressed files	10
compression	30, 57
console	
overview	23
requirements	5, 10
security	24
continuous data protection	See TimeData

D

dynamic volumes	10
-----------------	----

E

e-mail	
recovery server	58
repository server	31
source	99

encrypted files	10
error codes	117
event messages	68, 99
Exchange	See Ontrack PowerControls

F

file attributes	10
firewall	131

G

getting started	1
-----------------	---

I

import	23
installation	
automatic process	16
continuous data protection	19
installing Double-Take Backup	14
notes	13
Ontrack PowerControls	20
overview	13
recovery server	36
TimeData	19

J

job	
controls	60
import	23
monitoring	60
monitoring tools	60
optional protection settings	29
optional recovery settings	56
junction points	10

L

licensing	21
log file	67

M

Macintosh files	10
monitoring tools	
console	60
error codes	117
event messages	68
log files	67
overview	60
SNMP	111-112, 114
statistics	100-102
mount points	10

O

Ontrack PowerControls	
installing	20
overview	2-3
requirements	5, 8
viewing data	33, 35
overview	
archiving	2, 4
Cargo	2, 4
continuous data protection	2-3
Double-Take Backup	2
Ontrack PowerControls	2-3
TimeData	2-3

P

Performance Monitor	108
ports	131
preview report	124
protection	
establishing	25
optional settings	29
overview	25

R

recalling	129
-----------	-----

recovery	
data	53
full server to existing	37
full server to Hyper-V	48
full server to VMware ESX	41
installing to recovery server	36
optional settings	56
overview	36
recovery server	
defined	2
installation	36
requirements	5, 7
reparse points	10
replication capabilities	10
repository server	
defined	2
finding data	33
requirements	5-6
requirements	
archiving	9
Cargo	9
console	10
continuous data protection	8
Ontrack PowerControls	8
overview	5
recovery server	7
repository server	6
source server	5
TimeData	8
resources	1
route	30, 57
S	
security credentials	24
shutdown options	37, 41, 48
silent install	16
snapshot	
protection	25

recover data	53
recover full server to existing	37
recover full server to Hyper-V	48
recover full server to VMware ESX	41
viewing	33
SNMP	
configuration	112
overview	111
statistics	114
traps	112
source	
defined	2
requirements	5
SQL Server	
upgrade	13, 20
statistics	
file	101
output	102
overview	100
Performance Monitor	108
SNMP	114
symbolic links	10
system state	2
T	
TimeData	
installing	19
overview	2-3
protection	25
recover data	53
recover full server to existing	37
recover full server to Hyper-V	48
recover full server to VMware ESX	41
requirements	5, 8
viewing data	33
transactional NTFS operations (TxF)	10
transmission	
bandwidth	30, 57

compression	30, 57
route	30, 57
TxF (transactional NTFS operations)	10

U

upgrade	
notes	13
SQL Server	20
upgrading Double-Take Backup	14

W

welcome	1
---------	---