



Version 8.0.0

Double-Take DR for Windows User's Guide

Notices

Double-Take DR for Windows User's Guide Version 8.0, Monday, March 06, 2017

Check your service agreement to determine which updates and new releases you may be eligible for.

If you need maintenance renewal, an upgrade license key, or other sales assistance, contact your reseller/distributor or a Double-Take Software sales representative.

If you need technical assistance, you can contact CustomerCare. All basic configurations outlined in the online documentation will be supported through CustomerCare. Your technical support center is dependent on the reseller or distributor you purchased your product from and is identified on your service agreement. If you do not have access to this agreement, contact CustomerCare and they will direct you to the correct service provider. To contact CustomerCare, you will need your serial number and license key.

Assistance and support for advanced configurations may be referred to a Pre-Sales Systems Engineer or to Professional Services.

Man pages are installed and available on Double-Take Linux servers. These documents are bound by the same Double-Take Software license agreement as the software installation.

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

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

© 2017 Double-Take Software, Inc. All rights reserved.


Contents

Chapter 1 Double-Take DR overview	5
Getting Started	6
Chapter 2 Double-Take DR requirements	7
Mirroring and replication capabilities	13
Chapter 3 Double-Take Console	18
Double-Take Console requirements	20
Console options	21
Chapter 4 Managing servers	24
Adding servers	34
Providing server credentials	37
Viewing server details	38
Editing server properties	40
General server properties	41
Server licensing	42
Server setup properties	44
Double-Take queue	48
Source server properties	52
Target server properties	54
E-mail notification configuration	56
Script credentials	58
Log file properties	59
Verification log	61
Viewing server events	63
Viewing server logs	64
Managing snapshots	66
Snapshot states	67
Firewalls	70
Chapter 5 Protection	71
Creating a protection job	72
Managing and controlling protection jobs	87
Viewing protection job details	98
Validating a protection job	102
Editing a protection job	103
Viewing a protection job log	105
Chapter 6 Recovery	107
Creating a recovery job	108
Managing and controlling recovery jobs	130
Viewing recovery job details	140
Validating a recovery job	144
Viewing a recovery job log	145

Chapter 7 Simulating protection or recovery	147
Chapter 8 Managing images	148
Chapter 9 Special network configurations	149
Firewalls	150
IP and port forwarding	151

Chapter 1 Double-Take DR overview

Double-Take DR is a comprehensive data and full server recovery solution. It allows you to protect a production server, known as a source, by replicating an image of the source to another server, known as a target repository server. The image of the source can be data only or it can contain data and the source server's system state, which is the server's configured operating system and the installed applications. In the event the source should fail, you can use the image of the source, stored on the target repository server, to quickly and easily create a new source.



1. Source servers can be located locally or over a WAN.
2. As changes are made on the source, Double-Take mirroring and replication keeps the image of the source stored on the target repository server up-to-date. Double-Take replicates, in real-time, only the file changes rather than copying entire files, allowing you to more efficiently use resources.
3. If a source fails, you can use the image of the source from the target repository server to quickly and easily create a new physical or virtual source.

Getting Started

Before you begin, you should also review the *Double-Take DR requirements* on page 7 to make sure your environment meets the necessary requirements. When you are ready to begin, complete the following steps, in order.

1. Install Double-Take on your source server and target repository server.
2. If you are going to be running the Double-Take Console from a machine other than the source or target repository server, install the console.
3. License and, if necessary, activate your source and target repository servers.
4. Create your protection job. See *Protection* on page 71.



See the *Double-Take Installation, Licensing, and Activation* document for complete installation and licensing details.

Chapter 2 Double-Take DR requirements

Each type of server must meet certain requirements. Additionally, verify the machine where you are running the console meets the *Double-Take Console requirements* on page 20, and review the *Mirroring and replication capabilities* on page 13 to understand the type of data that Double-Take protects.

- **Source server**—The source server can be a physical or virtual server running any of the following operating systems.
 - Windows 2008 R2 Service Pack 1 or later
 - Windows 2012 or Windows 2012 R2
 - Server Core is supported with the following caveats.
 - Server Core 2008 R2 SP1, Server Core 2012, and Server Core 2012 R2 are supported. However, Server Core 2008 and Server Core 2008 R2 with no service packs are not supported.
 - If you are recovering a full server image of a Server Core server, your recovery server must be the same Server Core version.
 - DNS updates are not supported for Server Core servers.
- **Target repository server**—The target server can be any physical or virtual Windows 2008 R2 Service Pack 1 or later, Windows 2012, or Windows 2012 R2 operating system.
- **Recovery server**—The recovery server for a data protection job can be any operating system from the source server list above. The recovery server for a full server protection job can also be any operating system from the source server list above, however, keep in mind that a recovery server may be in that list but may not be suitable to stand-in for a source in the event of a source failure. See *Target compatibility* on page 11 for additional information regarding an appropriate recovery server for a full server protection job.
- **File system**—Double-Take supports the NTFS file system. FAT, FAT32, and ReFS are not supported. For detailed information on other file system capabilities, see *Mirroring and replication capabilities* on page 13.
- **System memory**—The minimum system memory on each server is 1 GB.
- **Disk space for program files**—This is the amount of disk space needed for the Double-Take program files. The amount depends on your operating system version and ranges from 350-500 MB.



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

Make sure you have additional disk space for Double-Take queuing, logging, and so on.

- **Server name**—Double-Take includes Unicode file system support, but your server name must still be in ASCII format. If you have the need to use a server's fully-qualified domain name, your server cannot start with a numeric character because that will be interpreted as an IP address. Additionally, all Double-Take servers must have a unique server name.



If you need to rename a server that already has a Double-Take license applied to it, you should deactivate that license before changing the server name. That includes rebuilding a server or changing the case (capitalization) of the server name (upper or lower case or any combination of case). If you have already rebuilt the server or changed the server name or case, you will have to perform a host-transfer to continue using that license. See the *Double-Take Installation, Licensing, and Activation* document for complete details.

- **Time**—The clock on your Double-Take servers must be within a few minutes of each other, relative to UTC. Large time skews (more than five minutes) will cause Double-Take errors.
- **Protocols and networking**—Your servers must meet the following protocol and networking requirements.
 - Your servers must have TCP/IP with static IP addressing.
 - IPv4 only configurations are supported, IPv4 and IPv6 are supported in combination, however IPv6 only configurations are not supported
 - By default, Double-Take is configured for IPv6 and IPv4 environments, but the Double-Take service will automatically check the server at service startup and modify the appropriate setting if the server is only configured for IPv4.
 - If you are using IPv6 on your servers, your clients must be run from an IPv6 capable machine.
 - In order to properly resolve IPv6 addresses to a hostname, a reverse lookup entry should be made in DNS.
 - If you are using Double-Take over a WAN and do not have DNS name resolution, you will need to add the host names to the local hosts file on each server running Double-Take.
- **NAT support**—Double-Take supports IP and port forwarding in NAT environments with the following caveats.
 - Only IPv4 is supported.
 - Only standalone servers are supported.
 - Make sure you have added your server to the Double-Take Console using the correct public or private IP address. The name or IP address you use to add a server to the console is dependent on where you are running the console. Specify the private IP address of any servers on the same side of the router as the console. Specify the public IP address of any servers on the other side of the router as the console.
 - DNS failover and updates will depend on your configuration
 - Only the source or target can be behind a router, not both.
 - The DNS server must be routable from the target
- **Reverse lookup zone**—If you are using a DNS reverse lookup zone, then it must be Active Directory integrated. Double-Take is unable to determine if this integration exists and therefore cannot warn you during job creation if it doesn't exist.
- **DNS updates**—Full server protection jobs allow you to failover Microsoft DNS records so the source server name resolves to the recovery server IP addresses at failover time. To be able to set up and failover Microsoft DNS records, your environment must meet the following requirements.

- The source and recovery servers must be in the same domain.
- The recovery server must have WMI/DCOM connectivity to any DNS server that you have configured to be updated.
- Each server's network adapter must have the DNS suffix defined, and the primary DNS suffix must be the same on the source and recovery server. You can set the DNS suffix in the network adapters advanced TCP/IP settings or you can set the DNS suffix on the computer name. See the documentation for your specific operating system for details on configuring the DNS suffix.
- If you are using a DNS reverse lookup zone, then the forward zone must be Active Directory integrated. Double-Take is unable to determine if this integration exists and therefore cannot warn you during job creation if it doesn't exist. The zone should be set for secure only updates to allow for DNS record locking.

DNS updates are not supported for Server Core servers.

- **Microsoft .NET Framework**—Microsoft .NET Framework version 4.5.1 is required.
- **Windows firewall**—If you have Windows firewall enabled on your servers, there are two requirements for the Windows firewall configuration.
 - The Double-Take installation program will automatically attempt to configure ports 6320, 6325, and 6326 for Double-Take. If you cancel this step, you will have to configure the ports manually.
 - If you are using the Double-Take Console to push installations out to your Windows servers, you will have to open firewall ports for 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. Additionally, you will need to open firewall ports for SMB (server message block) communications which uses ports 135-139 and port 445, and you will need to open File and Printer Sharing. As an alternative, you can disable the Windows firewall temporarily until the push installations are complete.

See *Firewalls* on page 150 for instructions on handling firewalls in your environment.

- **Clusters**—Cluster configurations are not supported on the source or target. Your servers must be in a standalone to standalone configuration.
- **Snapshots**—You can take and failover to Double-Take snapshots with protection and recovery jobs.

Double-Take uses the Microsoft Volume Shadow Copy service (VSS) for snapshot capabilities. To use this functionality, your servers must meet the following requirements.

- **Snapshot location**—Snapshots are taken and stored inside the image of the source on the target, so be sure that you select a volume that has enough space for the image of the source with its snapshots, as well as any other source images and their snapshots.
- **Double-Take installation location**—In order to enable Double-Take snapshots, Double-Take must be installed on the system drive. If Double-Take is not installed on the system drive, snapshots will be disabled when enabling protection.
- **Server IP address**—If you have specified an IP address as the source server name, but that IP address is not the server's primary IP address, you will have issues with snapshot functionality. If you need to use snapshots, use the source's primary IP address or its name.

- **Snapshot limitations**—Sometimes taking a snapshot may not be possible. For example, there may not be enough disk space to create and store the snapshot, or maybe the target is too low on memory. If a snapshot fails, an Event message and a Double-Take log message are both created and logged.

There are also limitations imposed by Microsoft Volume Shadow Copy that impact Double-Take snapshots. For example, different Double-Take job types create different snapshot types, either client-accessible or non-client-accessible. VSS only maintains 64 client-accessible snapshots, while it maintains 512 non-client-accessible snapshots. If the maximum number of snapshots exists and another one is taken, the oldest snapshot is deleted to make room for the new one.

Another example is that Double-Take snapshots must be created within one minute because Volume Shadow Copy snapshots must be created within one minute. If it takes longer than one minute to create the snapshot, the snapshot will be considered a failure.

You must also keep in mind that if you are using extended functionality provided by Volume Shadow Copy, you need to be aware of the impacts that functionality may have on Double-Take. For example, if you change the location where the shadow copies are stored and an error occurs, it may appear to be a Double-Take error when it is in fact a Volume Shadow Copy error. Be sure and review any events created by the VolSnap driver and check your Volume Shadow Copy documentation for details.

You can use Volume Shadow Copy for other uses outside Double-Take, for example Microsoft Backup uses it. Keep in mind though that the driver for Volume Shadow Copy is started before the driver for Double-Take. Therefore, if you use snapshots on your source and you revert any files on the source that are protected by your job, Double-Take will not be aware of the revert and the file change will not be replicated to the target. The file change will be mirrored to the target during the next mirroring process.

Volume Shadow Copy snapshots are associated with the volume they belong to. Since Double-Take mirrors and replicates the data on the volume and not the volume itself, snapshots taken on the source cannot be used on the target's volume. Therefore, snapshots taken on the source are not mirrored or replicated to the target.

Target compatibility



As you review the items in the following list, remember that in a recovery job, the target that was protecting your original source is now the source of the recovery job. The server you are recovering to is the target server.

- **Operating system version**—The source and target must have the same operating system. For example, you cannot have Windows 2008 on the source and Windows 2012 on the target. The two servers do not have to have the same level of service pack or hotfix. Windows 2008 and 2012 and their R2 releases are considered different operating systems. Therefore, you cannot have Windows 2008 or 2012 on the source and the corresponding R2 version on the target. The Windows edition (Standard, Enterprise, and so on) does not have to be the same.
 - **Operating system language**—Your servers must be running the same Windows localized version. For example, if your source is running an English language version of Windows, your target must also be running an English language version of Windows. If your source is running a Japanese language version of Windows, your target must also be running a Japanese language version of Windows. This applies to all localized languages.
 - **Source and target preparation**—Uninstall any applications or operating system features that are not needed from both your source and target. For example, unused language packs will slow down recovery since there are thousands of extra files that need to be examined. Ideally, your target should be as clean and simple a configuration as possible.
-



During recovery, Double-Take cannot add protocol stacks and specific installations to a target that does not already have them. For example, VPN stacks, communication stacks, and services such as Microsoft RRAS (Routing and Remote Access Service) must be pre-installed on the target.

- **Storage Server Edition**—If you are using Windows Storage Server Edition, you will need to check with your NAS vendor to verify if there are technical or license restrictions on recovering an image of a server to different hardware.
- **Server role**—The target cannot be a domain controller. Ideally, the target should not host any functionality (file server, application server, and so on) because the functionality will be removed when failover occurs.

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

- **Processors**—There are no limits on the number or speed of the processors, but the source and the target should have at least the same number of processors. If the target has fewer processors or slower speeds than the source, there will be performance impacts for the users after failover.
- **Memory**—The target memory should be within 25% (plus or minus) of the source. If the target has much less memory than the source, there will be performance impacts for the users after failover.

- **Network adapters**—You must map at least one NIC from the source to one NIC on the target. If you have NICs on the source that are not being used, it is best to disable them. If the source has more NICs than the target, some of the source NICs will not be mapped to the target. Therefore, the IP addresses associated with those NICs will not be available after failover. If there are more NICs on the target than the source, the additional NICs will still be available after failover and will retain their pre-failover network settings.
- **File system format**—The source and the target must have the NTFS file system format on each server. FAT and FAT32 are no longer supported.
- **Logical volumes**—There are no limits to the number of logical volumes, although you are bound by operating system limits. For each volume you are protecting on the source, the target must have a matching volume. For example, if you are protecting drives C: and D: on the source, the target cannot have drives D: and E:. In this case, the target must also have drives C: and D:. Additional target volumes are preserved and available after failover with all data still accessible.
- **System path**—The source and the target must have the same system path. The system path includes the location of the Windows files, Program Files, and Documents and Settings.
- **Double-Take path**—Double-Take must be installed on the same path (volume and full directory path) on the source and the target.
- **Disk space**—The target must have enough space to store the data from the source. This amount of disk space will depend on the applications and data files you are protecting. The more data you are protecting, the more disk space you will need. The target must also have enough space to store, process, and apply the source's system state data. The size of the system state will depend on the operating system and architecture.

A copy of the source system state data will be staged on the target boot volume in a folder called Staging-SSM. You can predict (approximately) how much space you will need in this staging folder by calculating the size of the following folders on your boot volume.

- Documents and Settings
- Program Files
- Program Files (x86)
- Program Data
- Windows
- Users
- Any other folders you manually select for staging

If the target's boot volume does not have enough space to accommodate the source data and the staging folder, the job will become stuck in a retrying state and will be unable to complete synchronization. You should also have approximately 2-3 GB or more additional space on the target boot volume (beyond your calculation above) to ensure there is enough space for failover processing.

The following are rough estimates for the free space needed for the staging folder for different operating systems.

- **Windows 2008 R2**—at least 10 GB
- **Windows 2012**—at least 14 GB
- **Windows 2012 R2**—at least 15 GB

These minimums are for a clean operating system installation. Operating system customizations, installed applications, and user data will increase the disk space requirement.

Mirroring and replication capabilities

For Windows source servers, Double-Take mirrors and replicates file and directory data stored on any NTFS Windows file system. Mirrored and replicated items also include Macintosh files, compressed files, NTFS attributes and ACLs (access control list), dynamic volumes, files with alternate data streams, sparse files, encrypted files, and reparse points. Files can be mirrored and replicated across mount points, although mount points are not created on the target.

Double-Take does not mirror or replicate items that are not stored on the file system, such as physical volume data and registry based data. Additionally, Double-Take does not mirror or replicate NTFS extended attributes, registry hive files, Windows or any system or driver pagefile, system metadata files (\$LogFile, \$Mft, \$Bitmap, \$Extend\UsnJrnl, \$Extend\Quota, and \$Extend\ObjId), hard links, or the Double-Take disk-based queue logs. The only exception to these exclusions is for the full server job types. If you are protecting your system state and data using full server protection, Double-Take will automatically gather and replicate all necessary system state data, including files for the operating system and applications.

Note the following replication caveats.

1. FAT and FAT32 are not supported.
2. ReFS is not supported.
3. You cannot replicate from or to a mapped drive.
4. If any directory or file contained in your job 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.
5. 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.
6. If you select files with alternate data streams, keep in mind the following.
 - a. Alternate data streams are not included in the job size calculation. Therefore, you may see the mirror process at 99-100% 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 file attributes and data comparison 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.
7. 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 can access the file.
 - d. When remirroring, the entire file is transmitted every time, regardless of the remirror settings.
 - e. Verification cannot check encrypted files because of the encryption. If remirror is selected, the entire encrypted file will be remirrored to the target. Independent of the remirror option, all encrypted files will be identified in the verification log.
 - f. Empty encrypted files will be mirrored to the target, but if you copy or create an empty encrypted file within the 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.
 - g. When you are replicating encrypted files, a temporary file is created on both the source and target servers. The temporary file is automatically created in the same directory as the Double-Take 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.
8. 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.
9. Double-Take supports transactional NTFS (TxF) write operations, with the exception of TxF SavePoints (intermediate rollback points).
- a. With transactional NTFS and Double-Take 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 mirror, the transacted view of the data on the source is used. This means the data on the target will be the same as the transacted view of the data on the source. If there are pending transactions, the Double-Take **Target Data State** will indicate **Transactions Pending**. As the pending transactions are committed or aborted, Double-Take mirrors any necessary changes to the target. Once all pending transactions are completed, the **Target Data State** will update to **OK**.
- If you see the pending transactions state, you can check the Double-Take log file for a list of files with pending transactions. As transactions are committed or aborted, the list is updated until all transactions are complete, and the **Target Data State** is **OK**.
- b. During replication, transactional operations will be processed on the target identically as they are on the source. If a transaction is committed on the source, it will be committed on the target. If a transaction is aborted on the source, it will be aborted on the target.
 - c. When cutover occurs any pending transactions on the target will be aborted.
10. Double-Take supports Windows symbolic links and junction points. A symbolic link is a link (pointer) to a directory or file. Junction points are links to directories and volumes.
- a. If the link and the file/directory/volume are both in your job, both the link and the file/directory/volume are mirrored and replicated to the target.

- b. If the link is in the job, but the file/directory/volume it points to is not, only the link is mirrored and replicated to the target. The file/directory/volume that the link points to is not mirrored or replicated to the target. A message is logged to the Double-Take log identifying this situation.
 - c. If the file/directory/volume is in the job, but the link pointing to it is not, only the file/directory/volume is mirrored and replicated to the target. The link pointing to the file/directory/volume is not mirrored or replicated to the target.
 - d. Junction points that are orphans (no counterpart on the source) will be processed for orphan files, however, the contents of a junction point (where it redirects you) will not be processed for orphan files.
11. If you have the Windows NtfsDisable8dot3NameCreation setting enabled on the source but disabled on the target, there is a potential that you could overwrite and lose data on the target because of the difference in how long file names will be associated with short files names on the two servers. This is only an issue if there are like named files in the same directory (for example, longfilename.doc and longfi~1.doc in the same directory). To avoid the potential for any data loss, the NtfsDisable8dot3NameCreation setting should be the same on both the source and target.
 12. Double-Take can replicate paths up to 32,760 characters, although each individual component (file or directory name) is limited to 259 characters. Paths longer than 32760 characters will be skipped and logged.
 13. If you rename the root folder of a job, Double-Take interprets this operation as a move from inside the job to outside the job. Therefore, since all of the files under that directory have been moved outside the job and are no longer a part of the job, those files will be deleted from the target replica copy. This, in essence, will delete all of your replicated data on the target. If you have to rename the root directory of your job, make sure that the job is not connected.
 14. Keep in mind the following caveats when including and excluding data for replication.
 - a. Do not exclude Microsoft Word temporary files from your job. When a user opens a Microsoft Word file, a temporary copy of the file is opened. When the user closes the file, the temporary file is renamed to the original file and the original file is deleted. Double-Take needs to replicate both the rename and the delete. If you have excluded the temporary files from your job, the rename operation will not be replicated, but the delete operation will be replicated. Therefore, you will have missing files on your target.
 - b. When Microsoft SQL Server databases are being replicated, you should always include the tempdb files, unless you can determine that they are not being used by any application. Some applications, such as PeopleSoft and BizTalk, write data to the tempdb file. You can, most likely, exclude temporary databases for other database applications, but you should consult the product documentation or other support resources before doing so.
 - c. Some applications create temporary files that are used to store information that may not be necessary to replicate. If user profiles and home directories are stored on a server and replicated, this could result in a significant amount of unnecessary data replication on large file servers. Additionally, the \Local Settings\Temporary Internet Files directory can easily reach a few thousand files and dozens of megabytes. When this is multiplied by a hundred users it can quickly add up to several gigabytes of data that do not need to be replicated.
 - d. Creating jobs that only contain one file may cause unexpected results. If you need to replicate just one file, add a second file to the job to ensure the data is replicated to the correct location. (The second file can be a zero byte file if desired.)
 15. Double-Take does not replicate the last access time if it is the only thing that has changed. Therefore, if you are performing incremental or differential backups on your target machine, you

need to make sure that your backup software is using an appropriate flag to identify what files have been updated since the last backup. You may want to use the last modified date on the file rather than the date of the last backup.

16. Keep in mind the following caveats when using anti-virus protection.
 - a. Virus protection software on the target should not scan replicated data. If the data is protected on the source, operations that clean, delete, or quarantine infected files will be replicated to the target by Double-Take. If the replicated data on the target must be scanned for viruses, configure the virus protection software on both the source and target to delete or quarantine infected files to a different directory that is not in the job. If the virus software denies access to the file because it is infected, Double-Take will continually attempt to commit operations to that file until it is successful, and will not commit any other data until it can write to that file.
 - b. You may want to set anti-virus exclusions on your source to improve replication performance. There are risks associated with making exclusions, so implement them carefully. For more information, see the Microsoft article [822158 Virus scanning recommendations for Enterprise computers that are running currently supported versions of Windows](#).
 - c. If you are using avast! anti-virus software, it must be installed in its default installation location if you want to protect your sever with a full server protection job. If it is not in its default installation directory, cutover will fail.
17. SQL Server 2005 or later may not initialize empty space when the database size increases due to the auto grow feature. Therefore, there is nothing for Double-Take to replicate when this empty space is created. When the empty space is populated with data, the data is replicated to the target. A verification report will report unsynchronized bytes between the source and target due to the empty space. Since the space is empty, the data on the source and target is identical. In the event of a failure, the SQL database will start without errors on the target.
18. If you have reparse points in your data set, Double-Take will replicate the tag, unless it is a known driver. If it is a known driver, for example Microsoft SIS, Double-Take will open the file allowing the reparse driver to execute the file. In this case, the entire file will be replicated to the target (meaning the file is no longer a reparse point on the target and has all the data).
19. Keep in mind if you have reparse points in your data set, the reparse driver cannot be loaded on the target during protection. You must load the reparse driver on the target after cutover in order to access the data. Additionally, you cannot have reparse points in your data set if you are using same server migration because the server is functioning as both a source and target.
20. If you are using the Microsoft Windows Update feature, keep in mind the following caveats.
 - a. Schedule your Windows Update outside the times when a mirroring operation (initial mirror or remirror) is running. Windows updates that occur during a mirror may cause data integrity issues on the target.
 - b. In some cases, Windows Update may perform an NTFS transactional rollback before displaying the dialog box to reboot the computer. This rollback will cause a mirror. If that mirror completes before the reboot, the reboot will trigger another mirror, unless you have configured Double-Take to only mirror changed files on reboot.
 - c. You must resolve any Windows Update incomplete operations or errors before cutover. (Check the windowsupdate.log file.) Also, do not cutover if the target is waiting on a Windows Update reboot. If cutover occurs before the required Windows Update reboot, the target may not operate properly or it may not boot. You could also get into a situation where the reboot repeats indefinitely. One possible workaround for the reboot loop


condition is to access a command prompt through the Windows Recovery Environment and delete the file `\Windows\winsxs\pending.xml` file. You may need to take ownership of the file to delete it. Contact technical support for assistance with this process or to evaluate other alternatives. Before you contact technical support, you should use the Microsoft System Update Readiness Tool as discussed in [Microsoft article 947821](#). This tool verifies and addresses many Windows Update problems.

21. If you are using Windows deduplication, keep in mind the following caveats.
 - a. Deduplicated data on the source will be expanded to its original size on the target when mirrored. Therefore, you must have enough space on the target for this expansion, even if you have deduplication enabled on the target.
 - b. If you have deduplicated data on the target, mirroring and replication (like any other write process) will create a new file or new blocks of data. Existing blocks of deduplicated data will remain as they were until the next garbage collection.
 - c. If you are protecting an entire server, you must have the deduplication feature installed on both the source and target. It can be enabled or disabled independently on the two servers, but it must at least be installed on both of the servers.
 - d. After cutover, the amount of disk space on the failed over server will be incorrect until you run the deduplication garbage collection which will synchronize the disk space statistics.
22. If you are using Windows storage pools on your source, you must create the storage pool on the target before cutover.
23. Replication is not case-sensitive. For example, if you rename the file `Test.txt` to `test.txt`, that change will not be replicated to the target. You will have to delete the file on the target and when it is remirrored, the new case of the file name will be used.
24. If you are using Windows 2008 R2, virtual hard disks can be mounted and dismounted reusing the same drive letter. However, once you have established a job, you cannot mount a different virtual hard disk to the same drive letter used in your job. This could cause errors, orphan files, or possibly data corruption. If you must change drive letters associated with a virtual hard disk, delete the job, change the mounting, and then re-create the job.

Chapter 3 Double-Take Console

After you have installed the console, you can launch it by selecting **Double-Take, Double-Take Console** from your **Programs, All Programs,** or **Apps,** depending on your operating system.

The Double-Take Console is used to protect and monitor your servers and jobs. Each time you open the Double-Take Console, you start at the **Servers** page which allows you to view, edit, add, remove, or manage the servers in your console. You can also create a new job from this page.



At the bottom of the Double-Take Console, you will see a status bar. At the right side, you will find links for **Jobs with warnings** and **Jobs with errors**. This lets you see quickly, no matter which page of the console you are on, if you have any jobs that need your attention. Select this link to go to the **Jobs** page, where the appropriate **Filter: Jobs with warnings** or **Filter: Jobs with errors** will automatically be applied.



The first time you start the console, you will see the getting started screen tips on the **Servers** page. These tips walk you through the basic steps of adding a server to your console, installing Double-Take on that server, and creating a job on that server. If you do not want to see the tips, close them. If you want to reopen the tips after you have closed them, select **Help, Show Getting Started Tips**.

You can manually check for Double-Take updates by selecting **Help, Check for Updates**.

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

Double-Take Console requirements

You must meet the following requirements for the Double-Take Console.

- **Operating system**—The Double-Take Console can be run from a Windows source or target. It can also be run from a physical or virtual machine running Windows 10, Windows 8, or Windows 7 Service Pack 1 or later.
- **Microsoft .NET Framework**—Microsoft .NET Framework version 4.5.1 is required.
- **Screen resolution**—For best results, use a 1024x768 or higher screen resolution.



The Double-Take installation prohibits the console from being installed on Server Core. Because Windows 2012 allows you to switch back and forth between Server Core and a full installation, you may have the console files available on Server Core, if you installed Double-Take while running in full operating system mode. In any case, you cannot run the Double-Take Console on Server Core.

Console options

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

- **Monitoring**—This section is used to determine how the console monitors your Double-Take servers.
 - **Monitoring interval**—Specifies how often, in seconds, the console refreshes the monitoring data. The servers will be polled at the specified interval for information to refresh the console.
 - **Automatic retry**—This option will have the console automatically retry server login credentials, after the specified retry interval, if the server login credentials are not accepted. Keep in mind the following caveats when using this option.
 - This is only for server credentials, not job credentials.
 - A set of credentials provided for or used by multiple servers will not be retried for the specified retry interval on any server if it fails on any of the servers using it.
 - Verify your environment's security policy when using this option. Check your policies for failed login lock outs and resets. For example, if your policy is to reset the failed login attempt count after 30 minutes, set this auto-retry option to the same or a slightly larger value as the 30 minute security policy to decrease the chance of a lockout.
 - Restarting the Double-Take Console will automatically initiate an immediate login.
 - Entering new credentials will initiate an immediate login using the new credentials.
 - **Retry on this interval**—If you have enabled the automatic retry, specify the length of time, in minutes, to retry the login.
- **Server Communication**—This section is used to determine how the console communicates with your Double-Take servers.
 - **Default port for XML web services protocol**—Specifies the port that the console will use when sending and receiving data to Double-Take servers. By default, the port is 6325. Changes to the console port will not take effect until the console is restarted.
 - **Default port for legacy protocol**—If you are using an older Double-Take version, you will need to use the legacy protocol port. This applies to Double-Take versions 5.1 or earlier.
- **Diagnostics**—This section assists with console troubleshooting.
 - **Export Diagnostic Data**—This button creates a raw data file that can be used for debugging errors in the Double-Take Console. Use this button as directed by technical support.
 - **View Log File**—This button opens the Double-Take Console log file. Use this button as directed by technical support. You can also select **View, View Console Log File** to open the Double-Take Console log file.
 - **View Data File**—This button opens the Double-Take Console data file. Use this button as directed by technical support. You can also select **View, View Console Data File** to open the Double-Take Console data file.
- **License Inventory**—This section controls if the console contains a license inventory. This

feature may not appear in your console if your service provider has restricted access to it.

- **Enable license inventory**—This option allows you to use this console to manage the Double-Take licenses assigned to your organization. When this option is enabled, the **License Inventory** page is also enabled.
- **Default Installation Options**—All of the fields under the **Default Installation Options** section are used by the push installation on the **Install** page. The values specified here will be the default options used for the push installation.
 - **Activate online after install completes**—Specify if you want to activate your Double-Take licenses at the end of the installation. The activation requires Internet access from the console machine or the machine you are installing to. Activation will be attempted from the console machine first and if that fails, it will be attempted from the machine you are installing to. If you choose not to have the installation activate your licenses, you will have to activate them through the console license inventory or the server's properties page.
 - **Location of install folders**—Specify the parent directory location where the installation files are located. The parent directory can be local on your console machine or a UNC path.
 - **Windows**—Specify the parent directory where the Windows installation file is located. The default location is where the Double-Take Console is installed, which is `\Program Files\Vision Solutions\Double-Take`. The console will automatically use the `\x64` subdirectory which is populated with the Windows installation files when you installed the console. If you want to use a different location, you must copy the `\x64` folder and its installation file to the different parent directory that you specify.
 - **Linux**—For Linux servers, you have two choices.
 - If you copied the Linux installation files from your download to your Double-Take Console installation location, you must make sure they are in a `\Linux` subdirectory under the parent directory you specified for **Location of install folders**. Copy the Linux `.deb` or `.rpm` files from your download to the `\Linux` subdirectory. Make sure you only have a single version of the Linux installation files in that location. The push installation cannot determine which version to install if there are multiple versions in the `\Linux` subdirectory.
 - If you have already deployed your Linux virtual recovery appliance, specify the UNC path to the installers share on the appliance. For example, if your appliance is called `DTAppliance`, use the path `\DTAppliance\installers` for the **Location of install folders**. The console will automatically use the installation files in the `\Linux` subdirectory of this share location. (This option is only available if you deployed an appliance from an `.ova` file. There are no installers available if you installed Double-Take on your appliance manually.)
- **Default Windows Installation Options**—All of the fields under the **Default Installation Options** section are used by the push installation on the **Install** page. The values specified here will be the default options used for the push installation.
 - **Temporary folder for installation package**—Specify a temporary location on the server where you are installing Double-Take where the installation files will be copied and run.
 - **Installation folder**—Specify the location where you want to install Double-Take on each server. This field is not used if you are upgrading an existing version of Double-Take. In that case, the existing installation folder will be used.
 - **Queue folder**—Specify the location where you want to store the Double-Take disk queue on each server.

- **Amount of system memory to use**—Specify the maximum amount of memory, in MB, that can be used for Double-Take processing.
 - **Minimum free disk space**—This is the minimum amount of disk space in the specified **Queue folder** that must be available at all times. This amount should be less than the amount of physical disk space minus the disk size specified for **Limit disk space for queue**.
 - **Do not use disk queue**—This option will disable disk queuing. When system memory has been exhausted, Double-Take will automatically begin the auto-disconnect process.
 - **Unlimited disk queue**—Double-Take will use an unlimited amount of disk space in the specified **Queue folder** for disk queuing, which will allow the queue usage to automatically expand whenever the available disk space expands. When the available disk space has been used, Double-Take will automatically begin the auto-disconnect process.
 - **Limit disk space for queue**—This option will allow you to specify a fixed amount of disk space, in MB, in the specified **Queue folder** that can be used for Double-Take disk queuing. When the disk space limit is reached, Double-Take will automatically begin the auto-disconnect process.
-



If the servers you are pushing to do not have a C drive, make sure you update the folder fields because the Double-Take Console will not validate that the fields are set to a volume that does not exist and the installation will not start.

- **Default Linux Installation Options**—All of the fields under the **Default Installation Options** section are used by the push installation on the **Install** page. The values specified here will be the default options used for the push installation.
 - **Temporary folder for installation package**—Specify a temporary location on the server where you are installing Double-Take where the installation files will be copied and run.

Chapter 4 Managing servers

To manage the servers in your console, select **Servers** from the toolbar. The **Servers** page is for server management and job creation.

- **Add and remove servers**—You can add servers to and remove servers from the console.
- **View and edit**—You can view server details and edit Double-Take server properties.
- **Create job**—You can create a protection, migration, or recovery job for a selected server.
- **Server organization**—You can organize the servers that are in your console into groups, allowing you to filter the servers you are viewing based on your organization.

Review the following sections to understand the information and controls available on the **Servers** page.



If you have uninstalled and reinstalled Double-Take on a server, you may see the server twice on the **Servers** page because the reinstall assigns a new unique identifier to the server. One of the servers (the original version) will show with the red X icon. You can safely remove that server from the console.

Left pane

You can expand or collapse the left pane by clicking on the **Server Highlights** heading. This pane allows you to organize your servers into folders. The servers displayed in the top right pane will change depending on the server group folder selected in the left pane. Every server in your console session is displayed when the **All Servers** group is selected. If you have created and populated server groups under **My Servers**, then only the servers in the selected group will be displayed in the right pane.

Between the main toolbar and the left pane is a smaller toolbar. These toolbar options control the server groups in the left pane.

Create New Server Group

Creates a new server group below the selected group

Rename Server Group

Allows you to rename the selected server group

Delete Server Group

Deletes the selected server group. This will not delete the servers in the group, only the group itself.

Overflow Chevron

Displays any toolbar buttons that are hidden from view when the window size is reduced.

Top right pane

The top pane displays high-level overview information about your servers. You can sort the data within a column in ascending and descending order. You can also move the columns to the left or right of each other to create your desired column order. The list below shows the columns in their default left to right order.

Column 1 (Blank)

The first blank column indicates the machine type.

-  Double-Take source or target server which could be a physical server, virtual machine, or a cluster node
-  Double-Take source or target server which is a Windows cluster
-  vCenter server
-  ESX server
-  Double-Take Reporting Service server
-  Offline server which means the console cannot communicate with this machine.
-  Any server icon with a red circle with white X overlay is an error which means the console can communicate with the machine, but it cannot communicate with Double-Take on it.

Column 2 (Blank)

The second blank column indicates the security level

-  Processing—The console is attempting to communicate with machine.
-  Administrator access—This level grants full control.
-  Monitor only access—This level grants monitoring privileges only.
-  No security access—This level does not allow monitoring or control.

Server

The name or IP address of the server. If you have specified a reserved IP address, it will be displayed in parenthesis.

Activity

There are many different **Activity** messages that keep you informed of the server activity. Most of the activity messages are informational and do not require any administrator interaction. If you see error messages, check the server details. See *Viewing server details* on page 38.

Version

The Double-Take product version information, if any.

Licensing Status

The status of the license, if any, on the server. If your license is expired, any jobs using that server will be in an error state. If you have multiple licenses, the status will indicate the license that requires the soonest action. For example, if you have a Double-Take Move license that expires in two days and a Double-Take Availability license that must be activated within 10 days, the status will be for the Double-Take Move license.

Product

The Double-Take products, if any, licensed for the server

Bottom right pane

The details displayed in the bottom pane provide additional information for the server highlighted in the top pane. You can expand or collapse the bottom pane by clicking on the **Server Highlights** heading.

Name

The name or IP address of the server.

Operating system

The operating system of the server. This field will not be displayed if the console cannot connect to Double-Take on the server.

Product

The Double-Take products, if any, licensed for the server

Version

The product version information, if any

Serial Number

The serial number associated with the Double-Take license

Toolbar

The following options are available on the main toolbar of the **Servers** page. Some options are only available for a single selected server and others are available for multiple selected servers.

Create a New Job

The available job creation choices depend on the Double-Take licenses applied to your server.

- **Protect**—If you are licensed for Double-Take Availability or Double-Take DR, use the **Protect** option to create a protection job for the selected server.
- **Recover**—If you are licensed for Double-Take DR and have selected a target repository server, use the **Recover** option to create a recovery job.
- **Migrate**—If you are licensed for Double-Take Move or certain Double-Take Availability licenses, use the **Migrate** option to create a migration job for the selected server.

Add Servers

Adds a new server. This button leaves the **Servers** page and opens the **Add Servers** page. See *Adding servers* on page 34.

View Server Details

Views detailed information about a server. This button leaves the **Servers** page and opens the **View Server Details** page. See *Viewing server details* on page 38.

Edit Server Properties

Edits the server's properties and options. This button leaves the **Servers** page and opens the **Edit Server Properties** page. See *Editing server properties* on page 40.

Remove Server

Removes the server from the console.

Provide Credentials

Changes the login credentials that the Double-Take Console use to authenticate to a server. This button opens the **Provide Credentials** dialog box where you can specify the new account information. See *Providing server credentials* on page 37. You will remain on the **Servers** page after updating the server credentials. If your jobs use the same credentials, make sure you also update the credentials for any active jobs on the server. See the *Managing and controlling* section for your specific job type.

If you are using domain credentials for your Double-Take servers and you change those credentials, you will continue to receive a Windows Security pop-up in the Double-Take Console, even if you enter correctly updated credentials in the pop-up. This is an unavoidable Windows WCF communication issue, and you must update the credentials for the Double-Take servers using **Provide Credentials** in order to terminate the repeated pop-ups.

Manage Group Assignments

Allows you to assign, move, and remove the selected server from specific server groups. This buttons opens the Manage Group Assignments dialog box where you can assign and unassign the server to specific server groups. The server will appear in server groups marked with a checkmark, and will not appear in groups without a checkmark. Servers assigned to a server group will automatically appear in parent server groups.

Install

Installs or upgrades Double-Take on the selected server. This button opens the **Install** page where you can specify installation options.

Uninstall

Uninstalls Double-Take on the selected server.

View Server Events

Views Windows event messages for a server. This button leaves the **Servers** page and opens the **View Server Events** page. See the *Reference Guide* for a complete list of Windows event messages.

View Server Logs

Views the Double-Take logs messages for a server. This button opens the **Logs** window. This separate window allows you to continue working in the Double-Take Console while monitoring log messages. You can open multiple logging windows for multiple servers. When the Double-Take Console is closed, all logging windows will automatically close.

Activate Online

Activates licenses and applies the activation keys to servers in one step. You must have Internet access for this process. You will not be able to activate a license that has already been activated.

Refresh

Refreshes the status of the selected servers.

Search

Allows you to search the product or server name for items in the list that match the criteria you have entered.

Overflow Chevron

Displays any toolbar buttons that are hidden from view when the window size is reduced.

Right-click menu

The following options are available on the right-click menu of the **Servers** page. Some options are only available for a single selected server and others are available for multiple selected servers.

Protect

If you are licensed for Double-Take Availability or Double-Take DR, use the **Protect** option to create a protection job for the selected server.

Recover

If you are licensed for Double-Take DR and have selected a target repository server, use the **Recover** option to create a recovery job.

Migrate

If you are licensed for Double-Take Move or certain Double-Take Availability licenses, use the **Migrate** option to create a migration job for the selected server.

View Server Details

Views detailed information about a server. This button leaves the **Servers** page and opens the **View Server Details** page. See *Viewing server details* on page 38.

Edit Server Properties

Edits the server's properties and options. This button leaves the **Servers** page and opens the **Edit Server Properties** page. See *Editing server properties* on page 40.

Remove Server

Removes the server from the console.

Provide Credentials

Changes the login credentials that the Double-Take Console use to authenticate to a server. This button opens the **Provide Credentials** dialog box where you can specify the new account information. See *Providing server credentials* on page 37. You will remain on the **Servers** page after updating the server credentials. If your jobs use the same credentials, make sure you also update the credentials for any active jobs on the server. See the *Managing and controlling* section for your specific job type.

If you are using domain credentials for your Double-Take servers and you change those credentials, you will continue to receive a Windows Security pop-up in the Double-Take Console, even if you enter correctly updated credentials in the pop-up. This is an unavoidable Windows WCF communication issue, and you must update the

credentials for the Double-Take servers using **Provide Credentials** in order to terminate the repeated pop-ups.

Manage Group Assignments

Allows you to assign, move, and remove the selected server from specific server groups. This button opens the Manage Group Assignments dialog box where you can assign and unassign the server to specific server groups. The server will appear in server groups marked with a checkmark, and will not appear in groups without a checkmark. Servers assigned to a server group will automatically appear in parent server groups.

Install

Installs or upgrades Double-Take on the selected server. This button opens the **Install** page where you can specify installation options.

Uninstall

Uninstalls Double-Take on the selected server.

Copy

Copies the information for the selected servers. You can then paste the server information as needed. Each server is pasted on a new line, with the server information being comma-separated.

Paste

Pastes a new-line separated list of servers into the console. Your copied list of servers must be entered on individual lines with only server names or IP addresses on each line.

View Server Events

Views Windows event messages for a server. This button leaves the **Servers** page and opens the **View Server Events** page. See the *Reference Guide* for a complete list of Windows event messages.

View Server Logs

Views the Double-Take logs messages for a server. This button opens the **Logs** window. This separate window allows you to continue working in the Double-Take Console while monitoring log messages. You can open multiple logging windows for multiple servers. When the Double-Take Console is closed, all logging windows will automatically close.

Activate Online

Activates licenses and applies the activation keys to servers in one step. You must have Internet access for this process. You will not be able to activate a license that has already been activated.

Gather Support Diagnostics

Executes the diagnostic DTInfo utility which collects configuration data for use when reporting problems to technical support. It gathers Double-Take log files; Double-Take and system settings; network configuration information such as IP, WINS, and DNS addresses; and other data which may be necessary for technical support to troubleshoot issues. You will be prompted for a location to save the resulting file which is created with the information gathered. Because this utility is gathering several pieces of information, across the network to your console machine, it may take several minutes to complete the information gathering and sending the resulting file to the console machine.

View Replication Service Details

Views the replication service details for a server. This button opens the **Replication service view** window. This separate window allows you to continue working in the Double-Take Console while monitoring the replication service details. You can open multiple **Replication service view** windows for multiple servers. When the Double-Take Console is closed, all **Replication service view** windows will automatically close. If you do not want to open separate windows, you can switch between servers that are in your Double-Take Console from within the **Replication service view** window. See the *Reference Guide* for a complete list of replication details.

Refresh

Refreshes the status of the selected servers.

Adding servers

The first time you start the console, the **Servers** page is empty. In order to protect and monitor your servers, you must insert your servers and/or appliances in the console.

Inserting servers manually

1. Select **Get Started** from the toolbar.
2. Select **Add servers** and click **Next**.
3. On the **Manual Entry** tab, specify the server information.
 - **Server**—This is the name or IP address of the server or appliance to be added to the console.



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

If you are using a NAT environment, make sure you add your server to the Double-Take Console using the correct public or private IP address. The name or IP address you use to add a server to the console is dependent on where you are running the console. Specify the private IP address of any servers on the same side of the router as the console. Specify the public IP address of any servers on the other side of the router as the console.

-
- **User name**—For a server, specify a user that is a member of the **Double-Take Admin** or **Double-Take Monitors** security group on the server.



If you are using domain credentials for your Double-Take servers and you change those credentials, you will continue to receive a Windows Security pop-up in the Double-Take Console, even if you enter correctly updated credentials in the pop-up. This is an unavoidable Windows WCF communication issue, and you must update the credentials for the Double-Take servers in the Double-Take Console **Servers** page in order to terminate the repeated pop-ups.

-
- **Password**—Specify the password associated with the **User name** you entered.
 - **Domain**—If you are working in a domain environment, specify the **Domain**.
 - **Management Service port**—If you want to change the port used by the Double-Take Management Service, disable **Use default port** and specify the port number you want to use. This option is useful in a NAT environment where the console needs to be able to communicate with the server using a specific port number. Use the public or private port depending on where the console is running in relation to the server you are adding.
4. After you have specified the server or appliance information, click **Add**.
 5. Repeat steps 3 and 4 for any other servers or appliances you want to add.
 6. If you need to remove servers or appliances from the list of **Servers to be added**, highlight a server and click **Remove**. You can also remove all of them with the **Remove All** button.

7. When your list of **Servers to be added** is complete, click **OK**.

Inserting servers through Active Directory discovery

You can insert servers using Active Directory discovery.

1. Select **Get Started** from the toolbar.
2. Select **Add servers** and click **Next**.
3. Select the **Automatic Discovery** tab.
4. Click **Discover** to search Active Directory for servers running Double-Take.
5. If you need to remove servers from the list of **Servers to be added**, highlight a server and click **Remove**. You can also remove all of them with the **Remove All** button.
6. When your list of **Servers to be added** is complete, click **OK**.
7. Because the Active Directory discovery uses pass-through authentication, you will need to update the credentials for each server from the **Servers** page, so that explicit credentials can be used when you go to create a job. Click **Provide Credentials** and provide credentials for a user that has privileges to that server and is a member of the Double-Take Admin security group.

Importing and exporting servers from a server and group configuration file

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

To export a server and group configuration file, select **File, Export Servers**. Specify a file name and click **Save**. After the configuration file is exported, you can import it to another console.

When you are importing a console server and group configuration file from another console, you will not lose or overwrite any servers that already exist in the console. For example, if you have server alpha in your console and you insert a server configuration file that contains servers alpha and beta, only the server beta will be inserted. Existing group names will not be merged, so you may see duplicate server groups that you will have to manually update as desired.

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

Providing server credentials

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



If you are using domain credentials for your Double-Take servers and you change those credentials, you will continue to receive a Windows Security pop-up in the Double-Take Console, even if you enter correctly updated credentials in the pop-up. This is an unavoidable Windows WCF communication issue, and you must update the credentials for the Double-Take servers in the Double-Take Console in order to terminate the repeated pop-ups.

Viewing server details

Highlight a server on the **Servers** page and click **View Server Details** from the toolbar. The **View Server Details** page allows you to view details about that particular server. The server details vary depending on the type of server or appliance you are viewing.

Server name

The name or IP address of the server. If you have specified a reserved IP address, it will be displayed in parenthesis.

Operating system

The server's operating system version

Roles

The role of this server in your Double-Take environment. In some cases, a server can have more than one role.

- **Engine Role**—Source or target server
- **Image Repository Role**—A target for a DR protection job or a source for a DR recovery job
- **Reporting Service**—Double-Take Reporting Service server

Status

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

Activity

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

Connected via

The IP address and port the server is using for communications. You will also see the Double-Take protocol being used to communicate with server. The protocol will be XML web services protocol (for servers running Double-Take version 5.2 or later) or Legacy protocol (for servers running version 5.1 or earlier).

Version

The product version information

Access

The security level granted to the specified user

User name

The user account used to access the server

Licensing

Licensing information for the server

Source jobs

A list of any jobs from this server. Double-clicking on a job in this list will automatically open the **View Job Details** page.

Target jobs

A list of any jobs to this server. Double-clicking on a job in this list will automatically open the **View Job Details** page.

Images

If this server is being used as a repository server (either as a target receiving data to an image for protection or as a source sending data from an image for recovery), the images on the server will be displayed.

Snapshots

If you have images available, highlight one and any snapshots associated with that image will be displayed. As you highlight a different image, the snapshots associated with that image will be displayed. If there are no snapshots associated with an image, the list will be blank.

Virtual Disks

If you have images available, highlight one and the virtual disks associated with that image will be displayed. As you highlight a different image, the virtual disks associated with that image will be displayed.


Editing server properties

Right-click a server on the **Servers** page and select **Edit server properties**. The **Edit Server Properties** page allows you to view and edit properties for that server. Click on a heading on the **Edit Server Properties** page to expand or collapse a section of properties.

- *General server properties* on page 41—Identifies the server
- *Server licensing* on page 42—Views, adds, and removes license keys
- *Server setup properties* on page 44—Indicates how the server will act on startup and shutdown
- *Double-Take queue* on page 48—Configures the Double-Take queues
- *Source server properties* on page 52—Configures the source server
- *Target server properties* on page 54—Configures the target server
- *E-mail notification configuration* on page 56—Configures e-mail notification
- *Script credentials* on page 58—Specifies credentials to be used when executing custom scripts during mirroring
- *Log file properties* on page 59—Configures log files

General server properties

The general server properties identify the server.



The screenshot shows a dialog box titled "General" with a close button in the top-left corner. It contains the following settings:

- Default address:** A dropdown menu with "(None)" selected.
- Port:** A dropdown menu with "6320" selected.
- Encrypt network data**
- Automatic retry**
- Retry on this interval (minutes):** A dropdown menu with "30" selected.

- **Default address**—On a server with multiple NICs, you can specify which address Double-Take traffic will use. It can also be used on servers with multiple IP addresses on a single NIC. If you change this setting, you must restart the Double-Take service for this change to take effect.
- **Port**—The server uses this port to send and receive commands and operations between Double-Take servers. If you change the port, you must stop and restart the Double-Take service.
- **Encrypt network data**—Use this option to encrypt your data before it is sent from the source to the target. Both the source and target must be encryption capable (version 7.0.1 or later), however this option only needs to be enabled on the source or target in order to encrypt data. Keep in mind that all jobs from a source with this option enabled or to a target with this option enabled will have the same encryption setting. Changing this option will cause jobs to auto-reconnect and possibly remirror.
- **Automatic retry**—This option will have the target server automatically retry server login credentials for a job, after the specified retry interval, if the server login credentials are not accepted. Keep in mind the following caveats when using this option.
 - Because server logins for a job are controlled by the target, this setting is only applicable to target servers.
 - This is only for server credentials, not job credentials.
 - Verify your environment's security policy when using this option. Check your policies for failed login lock outs and resets. For example, if your policy is to reset the failed login attempt count after 30 minutes, set this auto-retry option to the same or a slightly larger value as the 30 minute security policy to decrease the chance of a lockout.
- **Retry on this interval**—If you have enabled the automatic retry, specify the length of time, in minutes, to retry the login.

Server licensing

Licensing identifies your Double-Take license keys.



The fields and buttons in the **Licensing** section will vary depending on your Double-Take Console configuration and the type of license keys you are using.

Licensing

Add license keys and activation keys:

Choose...

Add

Current license keys:

	Product	Serial Number	Expiration Date
	Double-Take Availability for Windows...	7323	10/12/2014

Remove Copy

Changes to licensing are applied immediately.

Activation

This server has one or more licenses that require activation. To fully activate your software you can use the activate-online feature to complete the activation process using your internet connection. Alternatively using the link below, you may obtain activation keys from our activation website and add them above to activate your software.

Server information:

Copy

[Get an activation key from our website](#) Activate Online

- **Add license keys and activation keys**—Your license key or activation key is a 24 character, alpha-numeric key. You can change your license key without reinstalling, if your license changes. To add a license key or activation key, type in the key or click **Choose from inventory** and select a key from your console's license inventory. Then click **Add**.



The license inventory feature cannot be enabled if your service provider has restricted access to it.

- **Current license keys**—The server's current license key information is displayed. To remove a key, highlight it and click **Remove**. To copy a key, highlight it and click **Copy**. To replace a key,

enter a new key and click **Add**. If you are replacing an unexpired key with the same version and serial number, you should not have to reactivate it and any existing jobs will continue uninterrupted. If you are replacing an unexpired key with a new version or new serial number or replacing an expired key, you will have to reactivate and remirror.

- **Activation**—If your license key needs to be activated, you will see an additional **Activation** section at the bottom of the **Licensing** section. To activate your key, use one of the following procedures.
 - **Activate online**—If you have Internet access, you can activate your license and apply the activated license to the server in one step by selecting **Activate Online**.



You will not be able to activate a license that has already been activated.

- **Obtain activation key online, then activate**—If you have Internet access, click the hyperlink in the **Activation** section to take you to the web so that you can submit your activation information. Complete and submit the activation form, and you will receive an e-mail with the activation key. Activate your server by entering the activation key in the **Add license keys and activations keys** field and clicking **Add**.
- **Obtain activation key offline, then activate**—If you do not have Internet access, go to <https://activate.doubletake.com> from another machine that has Internet access. Complete and submit the activation form, and you will receive an e-mail with the activation key. Activate your server by entering the activation key in the **Add license keys and activations keys** field and clicking **Add**.


The activation key is specific to this server. It cannot be used on any other server. If the activation key and server do not match, Double-Take will not run.



If your Double-Take DR license keys needs to be activated, you will have 14 days to do so.

If you need to rename a server that already has a Double-Take license applied to it, you should deactivate that license before changing the server name. That includes rebuilding a server or changing the case (capitalization) of the server name (upper or lower case or any combination of case). If you have already rebuilt the server or changed the server name or case, you will have to perform a host-transfer to continue using that license.

Server setup properties



Server setup properties indicate how the server will act on startup and shutdown.

- **Log statistics automatically**—If enabled, Double-Take statistics logging will start automatically when Double-Take is started.
- **Enable task command processing**—Task command processing is a Double-Take feature that allows you to insert and run tasks at various points during the replication of data. Because the tasks are user-defined, you can achieve a wide variety of goals with this feature. For example, you might insert a task to create a snapshot or run a backup on the target after a certain segment of data from the source has been applied on the target. This allows you to coordinate a point-in-time backup with real-time replication. Enable this option to enable task command processing, however to insert your tasks, you must use the Double-Take scripting language. See the *Scripting Guide* for more information. If you disable this option on a source server, you can still submit tasks to be processed on a target, although task command processing must be enabled on the target.
- **Advertise service with Active Directory**—If enabled, the Double-Take service registers with Windows Active Directory when the service is started.
- **Mirror only changed files when source reboots**—If enabled, Double-Take will use the Double-Take driver change journal and the Windows change journal to track file changes. If the source is rebooted, only the files identified in the change journals will be remirrored to the target. This setting helps improve mirror times. If this option is enabled but the change journals cannot be used or if this option is disabled, the selected choice for Behavior when automatically remirroring will be used to remirror changes after the source reboots.



If you reboot your source into safe mode and changes are made to the protected data and then the source is rebooted normally, the Double-Take driver change journal will try but not be able to synchronize the source and target correctly because it was not loaded in safe mode. Therefore, you should manually start a difference mirror.

- **Automatically reconnect during source initialization**—Disk queues are user configurable and can be extensive, but they are limited. If the amount of disk space specified for disk queuing is met, additional data would not be added to the queue and data would be lost. To avoid any data loss, Double-Take will automatically disconnect jobs when necessary. If this option is enabled, Double-Take will automatically reconnect any jobs that it automatically disconnected. These processes are called auto-disconnect and auto-reconnect and can happen in the following scenarios.

- **Source server restart**—If your source server is restarted, Double-Take will automatically reconnect any jobs that were previously connected. Then, if configured, Double-Take will automatically remirror the data. This process is called auto-remirror. The remirror re-establishes the target baseline to ensure data integrity, so disabling auto-remirror is not advised.
- **Exhausted queues on the source**—If disk queuing is exhausted on the source, Double-Take will automatically start disconnecting jobs. This is called auto-disconnect. The transaction logs and system memory are flushed allowing Double-Take to begin processing anew. The auto-reconnect process ensures that any jobs that were auto-disconnected are automatically reconnected. Then, if configured, Double-Take will automatically remirror the data. This process is called auto-remirror. The remirror re-establishes the target baseline to ensure data integrity, so disabling auto-remirror is not advised.
- **Exhausted queues on the target**—If disk queuing is exhausted on the target, the target instructs the source to pause. The source will automatically stop transmitting data to the target and will queue the data changes. When the target recovers, it will automatically tell the source to resume sending data. If the target does not recover by the time the source queues are exhausted, the source will auto-disconnect as described above. The transaction logs and system memory from the source will be flushed then Double-Take will auto-reconnect. If configured, Double-Take will auto-remirror. The remirror re-establishes the target baseline to ensure data integrity, so disabling auto-remirror is not advised.
- **Queuing errors**—If there are errors during disk queuing on either the source or target, for example, Double-Take cannot read from or write to the transaction log file, the data integrity cannot be guaranteed. To prevent any loss of data, the source will auto-disconnect and auto-reconnect. If configured, Double-Take will auto-remirror. The remirror re-establishes the target baseline to ensure data integrity, so disabling auto-remirror is not advised.
- **Target server interruption**—If a target machine experiences an interruption (such as a cable or NIC failure), the source/target network connection is physically broken but both the source and target maintain the connection information. The Double-Take source, not being able to communicate with the Double-Take target, stops transmitting data to the target and queues the data changes, similar to the exhausted target queues described above. When the interruption is resolved and the physical source/target connection is reestablished, the source begins sending the queued data to the target. If the source/target connection is not reestablished by the time the source queues are exhausted, the source will auto-disconnect as described above.
- **Target service shutdown**—If the target service is stopped and restarted, there could have been data in the target queue when the service was stopped. To prevent any loss of data, the Double-Take service will attempt to persist to disk important target connection information (such as the source and target IP addresses for the connection, various target queue information, the last acknowledged operation, data in memory moved to disk, and so on) before the service is stopped. If Double-Take is able to successfully persist this information, when the Double-Take service on the target is restarted, Double-Take will pick up where it left off, without requiring an auto-disconnect, auto-reconnect, or auto-remirror. If Double-Take cannot successfully persist this information prior to the restart (for example, a server crash or power failure where the target service cannot shutdown gracefully), the source will auto-reconnect when the target is available, and if configured, Double-Take will auto-remirror. The remirror re-establishes the target baseline to ensure data integrity, so disabling auto-remirror is not advised.



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

If you have manually changed data on the target, for example if you were testing data on the target, Double-Take is unaware of the target data changes. You must manually remirror your data from the source to the target, overwriting the target data changes that you caused, to ensure data integrity between your source and target.

- **Behavior when automatically remirroring**—Specify how Double-Take will perform the mirror when it is automatically remirroring.
-



If you are using files and folders, full server to ESX appliance, files and folders migration, DR data protection, or DR data recovery job and are using a database application or are protecting a domain controller, do not use the compare file attributes only options unless you know for certain that you need it. With database applications and because domain controllers store their data in a database, it is critical that all files, not just some of the files, are mirrored. In this case, you should compare both the attributes and the data.

- **Do not compare files. Send the entire file.**—Double-Take will not perform any comparisons between the files on the source and target. All files will be mirrored to the target, sending the entire file. This is equivalent to selecting the mirror all files option prior to Double-Take version 7.1.
- **Compare file attributes. Send the entire file.**—Double-Take will compare file attributes and will mirror those files that have different attributes, sending the entire file. This is equivalent to selecting the mirror different files and the only if source is newer options prior to Double-Take version 7.1. This option is not available for the following jobs.
 - Full server
 - Full server for Linux
 - SQL
 - Full server (Windows) to ESX
 - Full server (Linux) to ESX
 - Full server to Hyper-V
 - Full server migration
 - Full server to ESX migration
 - Full server to Hyper-V migration
 - Full server to Windows DR protection
 - Full server to Windows DR recovery
- **Compare file attributes. Send the attributes and bytes that differ.**—Double-Take will compare file attributes and will mirror only the attributes and bytes that are different. This is equivalent to selecting the mirror different files, only if source is newer, and use block checksum options prior to Double-Take version 7.1. This option is not available for the following jobs.


- Full server
- Full server for Linux
- SQL
- Full server (Windows) to ESX
- Full server (Linux) to ESX
- Full server to Hyper-V
- Full server migration
- Full server to ESX migration
- Full server to Hyper-V migration
- Full server to Windows DR protection
- Full server to Windows DR recovery
- **Compare file attributes and data. Send the attributes and bytes that differ.**—Double-Take will compare file attributes and the file data and will mirror only the attributes and bytes that are different. This is equivalent to selecting the mirror different files and use block checksum options prior to Double-Take version 7.1. If you are using a database application on your source, select this option.
- **Do not mirror**—Double-Take will not perform any comparisons between the files on the source and target. No files will be mirrored to the target.
- **Time allowed to complete shutdown operations**—This setting indicates the amount of time, in seconds, for the Double-Take service to wait prior to completing a shutdown so that Double-Take can persist data on the target in an attempt to avoid a remirror when the target comes back online. A timeout of zero (0) indicates waiting indefinitely and any other number indicates the number of seconds. The timeout setting only controls the service shutdown caused by Double-Take. It does not control the service shutdown through a reboot or from the Service Control Manager.

Double-Take queue

During the Double-Take installation, you identified the amount of disk space that can be used for Double-Take queuing. Queuing to disk allows Double-Take to accommodate high volume processing that might otherwise exhaust system memory. For example, on the source, this may occur if the data is changing faster than it can be transmitted to the target, or on the target, a locked file might cause processing to back up.

Double-Take Queuing Diagram

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



1. If data cannot immediately be transmitted to the target, it is stored in system memory. You can configure how much system memory you want Double-Take to use for all of its processing.
2. When the allocated amount of system memory is full, new changed data bypasses the full system memory and is queued directly to disk. Data queued to disk is written to a transaction log. Each transaction log can store 5 MB worth of data. Once the log file limit has been reached, a new transaction log is created. The logs can be distinguished by the file name which includes the target IP address, the Double-Take port, the connection ID, and an incrementing sequence number.



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

3. When system memory is full, the most recent changed data is added to the disk queue, as described in step 2. This means that system memory contains the oldest data. Therefore, when data is transmitted to the target, Double-Take pulls the data from system memory and sends it. This ensures that the data is transmitted to the target in the same order it was changed on the source. Double-Take automatically reads operations from the oldest transaction log file into

system memory. As a transaction log is depleted, it is deleted. When all of the transaction log files are deleted, data is again written directly to system memory (step 1).

4. To ensure the integrity of the data on the target, the information must be applied in the same order as it was on the source. If there are any delays in processing, for example because of a locked file, a similar queuing process occurs on the target. Data that cannot immediately be applied is stored in system memory.
5. When the allocated amount of system memory on the target is full, new incoming data bypasses the full system memory and is queued directly to disk. Data queued to disk is written to a transaction log. On the target, the transaction logs are identified with the source IP address, the Double-Take port, the connection ID, and an incrementing sequence number.

Like the source, system memory on the target contains the oldest data so when data is applied to the target, Double-Take pulls the data from system memory. Double-Take automatically moves operations from the oldest transaction log file to system memory. As a transaction log is depleted, it is deleted. When all of the transaction log files are deleted, data is again written directly to system memory (step 4).

The following memory and queue options are available for each Double-Take server.

The screenshot shows the 'Queue' configuration window. It includes a 'Queue folder' field with the path 'C:\Program Files\Vision Solutions\Double-Take\'. Below this is a note: 'A dedicated non-boot volume will offer the best performance and reliability.' There are two spinners: 'Amount of system memory to use (MB)' set to 1023 and 'Minimum free disk space (MB)' set to 250. There are three radio buttons: 'Do not use disk queue' (unselected), 'Unlimited disk queue' (selected), and 'Limit disk space for queue (MB)' (unselected). The 'Alert at this queue usage (%)' spinner is set to 50. At the bottom, there is a note: 'Changing the queue location or memory usage will not take effect until the service has been restarted.'

- **Queue folder**—This is the location where the disk queue will be stored. Any changes made to the queue location will not take effect until the Double-Take service has been restarted on the server.

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

- Select a location on a non-clustered volume that will have minimal impact on the operating system and applications.
- Select a location that is on a different volume as the location of the Windows pagefile.
- Select a dedicated, non-boot volume.
- Do not select the root of a volume.
- Do not select the same physical or logical volume as the data being replicated.
- On a Windows 2012 server, do not select a volume where deduplication is enabled.

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

mirrored or non-parity striped set because writing to a RAID 5 array incurs the overhead of generating and writing parity data. RAID 5 write performance can be up to 50% less than the write performance of a single disk, depending on the adapter and disk.



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

- **Amount of system memory to use**—This is the maximum amount of Windows system memory, in MB, that Double-Take will use. When this limit is reached, queuing to disk will be triggered. The minimum amount of system memory is 512 MB. The maximum amount is dependent on the server hardware and operating system. If you set this value lower, Double-Take will use less system memory, but you will queue to disk sooner which may impact system performance. If you set it higher, Double-Take will maximize system performance by not queuing to disk as soon, but the system may have to swap the memory to disk if the system memory is not available.

Since the source is typically running a production application, it is important that the amount of memory Double-Take and the other applications use does not exceed the amount of RAM in the system. If the applications are configured to use more memory than there is RAM, the system will begin to swap pages of memory to disk and the system performance will degrade. For example, by default an application may be configured to use all of the available system memory when needed, and this may happen during high-load operations. These high-load operations cause Double-Take to need memory to queue the data being changed by the application. In this case, you would need to configure the applications so that they collectively do not exceed the amount of RAM on the server. Perhaps on a server with 4 GB of RAM running the application and Double-Take, you might configure the application to use 1 GB and Double-Take to use 1 GB, leaving 2 GB for the operating system and other applications on the system. Many server applications default to using all available system memory, so it is important to check and configure applications appropriately, particularly on high-capacity servers.

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

- **Do not use disk queue**—This option will disable disk queuing. When system memory has been exhausted, Double-Take will automatically begin the auto-disconnect process.
- **Unlimited disk queue**—Double-Take will use an unlimited amount of disk space in the specified **Queue folder** for disk queuing, which will allow the queue usage to automatically expand whenever the available disk space expands. When the available disk space has been used, Double-Take will automatically begin the auto-disconnect process.
- **Limit disk space for queue**—This option will allow you to specify a fixed amount of disk space, in MB, in the specified **Queue folder** that can be used for Double-Take disk queuing. When the disk space limit is reached, Double-Take will automatically begin the auto-disconnect process.
- **Minimum free disk space**—This is the minimum amount of disk space in the specified **Queue folder** that must be available at all times. This amount should be less than the amount of physical disk space minus the disk size specified for **Limit disk space for queue**.



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

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

Source server properties

These properties are specific to the source server role.

Source

Number of replication packets per one mirror packet:
5
Changing this ratio does not affect current connections.

Replicate NTFS security attributes by name

Maximum pending mirror operations:
1000

Size of mirror packets (bytes):
65536

Calculate size of protected data upon connection

- **Number of replication packets per one mirror packet**—You can specify the ratio of replication packets to mirror packets that are placed in the source queue. The default value (5) allows Double-Take to dynamically change the ratio as needed based on the amount of replication data in queue. If you set a specific value other than the default (other than 5), the specified value will be used. Changes to this setting will take effect for future jobs. Existing jobs will have to be stopped and restarted to pick up the new ratio.
- **Replicate NTFS security attributes by name**—If you are protecting or migrating data, Double-Take allows you to replicate Windows permission attributes by local name as well as security ID (SID). By replicating Windows security by name, you can transmit the owner name with the file. If that user exists on the target, then the SID associated with the user will be applied to the target file ownership. If that user does not exist on the target, then the ownership will be unknown. By default, this option is disabled.
 - **Domain security model**—If you are using a Windows domain security model by assigning users at the domain level, each user is assigned a security ID (SID) at the domain level. When Double-Take replicates a file to the target, the SID is also replicated. Since a user will have the same SID on the source and target, the user will be able to access the file from the target. Therefore, this option is not necessary.
 - **Local security model**—If you are using a Windows local security model by assigning users at the local level (users that appear on multiple machine will each have different SIDs), you will need to enable this feature so that users can access the data on the target. If you do not enable this feature with a local security model, after a Double-Take file and SID is replicated, a local user will not be able to access the file because the user's SID on the target machine is different from the SID that was replicated from the source machine.

If you enable this option, make sure that the same groups and users exist on the target as they do on the source.

Enabling this option may have an impact on the rate at which Double-Take can commit data on the target. File security attributes are sent to the target during mirroring and replication. The target must obtain the security ID (SID) for the users and groups that are assigned permissions, which takes some time. If the users and groups are not on the target server, the delay can be substantial. The performance impact of enabling this option will vary depending on the type of file activity and other variables. For instance, it will not affect the overall performance of large database files much

(since there is a lot of data, but only a few file permissions), but may affect the performance of user files significantly (since there are often thousands of files, each with permissions). In general, the performance impact will only be noticed during mirrors since that is when the target workload is greatest.

Regardless of the security model you are using, if you create new user accounts on the source, you should start a remirror so the new user account information associated with any files in your job can be transmitted to the target.

- **Maximum pending mirror operations**—This option is the maximum number of mirror operations that are queued on the source. The default setting is 1000. If, during mirroring, the mirror queued statistic regularly shows low numbers, for example, less than 50, this value can be increased to allow Double-Take to queue more data for transfer.
- **Size of mirror packets**—This option determines the size of the mirror packets, in bytes, that Double-Take transmits. The default setting is 65536 bytes. You may want to consider increasing this value in a high latency environment (greater than 100 ms response times), or if your data set contains mainly larger files, like databases.
- **Calculate size of protected data upon connection**—Specify if you want Double-Take to determine the mirroring percentage calculation based on the amount of data being protected. If you enable this option, the calculation will begin when mirroring begins. For the initial mirror, the percentage will display after the calculation is complete, adjusting to the amount of the mirror that has completed during the time it took to complete the calculation. Subsequent mirrors will initially use the last calculated size and display an approximate percentage. Once the calculation is complete, the percentage will automatically adjust down or up to indicate the amount that has been completed. Disabling calculation will result in the mirror status not showing the percentage complete or the number of bytes remaining to be mirrored.



The calculated amount of protected data may be slightly off if your data set contains compressed or sparse files.


Do not disable this option for any of the following job types.

- Full server jobs
- Full server to Hyper-V jobs
- Full server to ESX jobs

The calculation time is when the source generates a hard link report that the target will use to process hard links on the target. If you disable the calculation and thus the report generation, the hard link processing that occurs on the target will be ineffective and you may have problems after failover.

Target server properties

These properties are specific to the target server role.



- **Pause mirroring at this level**—You can specify the maximum percentage of Windows system memory that can contain mirror data before the target signals the source to pause the sending of mirror operations. The default setting is 20.
- **Resume mirroring at this level**—You can specify the minimum percentage of Windows system memory that can contain mirror data before the target signals the source to resume the sending of mirror operations. The default setting is 15. You cannot set the resume value higher than the pause value.
- **Retry delay for incomplete operations**—This option specifies the amount of time, in seconds, before retrying a failed operation on the target. The default setting is 3.
- **Block target paths on connection**—You can block writing to the replica source data located on the target. This keeps the data from being changed outside of Double-Take processing.



Do not block your target paths if you are protecting an entire server because system state data will not be able to be written to the target.

Be careful blocking target paths if you will be using Double-Take snapshots. You will have to unblock the paths before you can failover to a snapshot. Additionally, be careful when blocking target paths with backup software running on the target. You will need to unblock the paths to allow backup software to take snapshots or update archive bits.

- **Move deleted files to this folder**—This option allows you to save files that have been deleted, by moving them to a different location on the target. When a file deletion is replicated to the target, instead of the file being deleted from the target, the file is moved to the specified location. This allows for easy recovery of those files, if needed. If you enable this option, specify where you want to store the deleted files.



If you are moving deleted files on the target and you have orphan files configured for removal (which is the default setting for most job types), do not move the deleted files to a

location inside the replica data on the target. The deleted files that are moved will then be deleted by the orphan file functionality.

- **Remove deleted files after this number of days**—If you are moving deleted files, you can specify a length of time, in days, to maintain the moved files. A moved file that is older than the specified number of days will be deleted. Double-Take checks for moved files that should be deleted once daily at 8 PM. Only the date, not the time, of the file is considered when moved files are deleted. For example, if you specify to delete moved files after 30 days, any file that is 31 days old will be deleted. Because the criteria is based on days and not time, a file that will be deleted could have been moved anytime between 12:01 AM and 11:59 PM 31 days ago.
-



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

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


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

Encrypted files that are deleted on the source will only be moved on the target if the move location is on the same volume as the copy of the source data on the target.

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

E-mail notification configuration

You can email Double-Take event messages to specific addresses using an SMTP mail server. 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 event message will be displayed in the body of the e-mail message.



E-mail Notification

Enable e-mail notifications Test

E-mail server (SMTP):
MailServer

Log on to e-mail server

User name:
administrator

Password:
●●●●●●●●

From address:
admin@domain.com

Send to:
admin@domain.com, ITSupport@domain.com

Subject prefix:
Double-Take Notification

Add event description to subject

Include these events:

Information
 Warning
 Error

Exclude these event IDs:
[Empty text box]

Enter event IDs and event ID ranges separated by commas. For example: 4000,4002,4010-4020.

- **Enable e-mail notification**—This option enables the e-mail notification feature. Any specified notification settings will be retained if this option is disabled.
- **E-mail server**—Specify the name of your SMTP mail server.
- **Log on to e-mail server**—If your SMTP server requires authentication, enable this option and specify the **User name** and **Password** to be used for authentication. Your SMTP server must support the LOGIN authentication method to use this feature. If your server supports a different authentication method or does not support authentication, you may need to add the Double-Take server as an authorized host for relaying e-mail messages. This option is not necessary if you are sending exclusively to e-mail addresses that the SMTP server is responsible for.
- **From address**—Specify the e-mail address that you want to appear in the From field of each Double-Take e-mail message. The address is limited to 256 characters.

- **Send to**—Specify the e-mail addresses that each Double-Take e-mail message should be sent to. Enter the addresses as a comma or semicolon separated list. Each address is limited to 256 characters. You can add up to 256 e-mail addresses.
- **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 100 characters.

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

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

- **Includes these events**—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.



When you modify your e-mail notification settings, you will receive a test e-mail summarizing your new settings. You can also test e-mail notification by clicking **Test**. By default, the test will be run from the machine where the 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.

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


If an error occurs while sending an e-mail, a message will be generated. This message will not trigger another 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 anti-virus software blocks unknown processes from sending traffic on port 25. You need to modify the blocking rule so that Double-Take e-mail messages are not blocked.

Script credentials

These credentials will be used when executing custom scripts for mirroring.



The screenshot shows a dialog box titled "Script Credentials" with a close button in the top-left corner. Below the title bar, there is a descriptive text: "These credentials will be used when executing custom scripts, such as during mirroring or failover." The dialog contains three input fields: "User name:" with the text "administrator", "Password:" with a masked password of ten dots, and "Domain:" with the text "domain.com". At the bottom of the dialog is a "Test" button.

Specify a **User name**, **Password**, and **Domain** to use when running the scripts. If you do not specify any security credentials, the account running the Double-Take service will be used. After you have specified credentials, you can click **Test** to confirm the credentials can be used for a successful login. If the credentials cannot be authenticated, you will receive an error. You will need to manually test that credentials you supply have appropriate rights to execute any scripts you may be running.

Log file properties

These settings allow you to specify your log file configuration.

The screenshot shows a dialog box titled "Logging" with several sections:

- Logging folder:** A text box containing "C:\Program Files\Vision Solutions\Double-Take\" and a browse button (...).
- Messages & Alerts:** Two spinners: "Maximum size (bytes)" set to 5242880 and "Maximum number of files" set to 20.
- Verification:** A text box for "File name" containing "DTVerify.log", a spinner for "Maximum size (bytes)" set to 1048576, a checked "Append" checkbox, and a dropdown for "Language" set to "English (United States)".
- Statistics:** A text box for "File name" containing "statistic.sts", a spinner for "Maximum size (bytes)" set to 10485760, and a spinner for "Write interval (minutes)" set to 5.

- **Logging folder**—Specify the directory where each of the log files in this section are stored. The default location is the directory where the Double-Take program files are installed.
- **Messages & Alerts**—These settings apply to the service log file.
 - **Maximum size**—Specify the maximum size, in bytes, of the log file. The default size is 5242880 bytes (5 MB). Once the maximum has been reached, a new log file will be created.
 - **Maximum number of files**—Specify the maximum number of log files that are maintained. The default is 20, and the maximum is 999. Once the maximum has been reached, the oldest file will be overwritten.
- **Verification**—The verification log is created during the verification process and details which files were verified as well as the files that are synchronized. See *Verification log* on page 61.
 - **File name**—This field contains the base log file name for the verification process. The job type and a unique identifier will be prefixed to the base log file name. For example, since the default is DTVerify.log, the verification log for a files and folders job will be Files and Folders_123456abcdef DTVerify.log.
 - **Maximum size**—Specify the maximum size, in bytes, of the verification log file. The default is 1048576 bytes (1 MB).
 - **Append**—Enable the **Append** check box if you want to append each verification process to the same log file. If this check box is disabled, each verification process that is logged will overwrite the previous log file. By default, this option is enabled.
 - **Language**—Select the language for your verification log file.

- **Statistics**—The statistics log maintains connection statistics such as mirror bytes in queue or replication bytes sent. This file is a binary file that is read by the DTStat utility. See the *Reference Guide* for details on DTStat.
 - **File name**—This is the name of the statistics log file. The default file name is statistic.sts.
 - **Maximum size**—Specify the maximum size, in bytes, of the statistics log file. The default is 10485760 bytes (10 MB). Once this maximum has been reached, the oldest data will be overwritten.
 - **Write interval**—Specify how often, in minutes, Double-Take writes to the statistics log file. The default is every 5 minutes.

Verification log

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

Sample verification log

```
--- VERIFICATION OF CONNECTION 2, CHECKSUM ENABLED (Sales data for alpha --> 206.31.65.40 : 1100) ---
Start Time: 1/24/2017 12:15:20 PM for connection 2 (Sales data for alpha -->
206.31.65.40 : 1100)
File:      beta\users\bob\budget.xls DIFFERENT ON TARGET
  Source Attributes: Timestamp = 1/17/2017 8:21:36 PM Size = 1272 Mask = [0x20]
  Target Attributes: Timestamp = 1/17/2017 8:21:36 PM Size = 1272 Mask = [0x20]
Security descriptors are different.
  0 BYTES OUT OF SYNC
File:      beta\users\bill\timesheet.xls DIFFERENT ON TARGET
  Source Attributes: Timestamp = 1/17/2017 8:21:37 PM Size = 1272 Mask = [0x20]
  Target Attributes: Timestamp = 1/17/2017 8:21:37 PM Size = 1272 Mask = [0x23]
  0 BYTES OUT OF SYNC
File:      beta\users\vincent\training.doc DIFFERENT ON TARGET
  Source Attributes: Timestamp = 1/12/2017 3:28:20 PM Size = 17 Mask = [0x20]
  Target Attributes: Timestamp = 1/20/2017 5:05:26 PM Size = 2 Mask = [0x20]
  17 BYTES OUT OF SYNC
Completion Time: 1/24/2017 12:37:44 PM for connection 2 (Sales data for alpha -->
206.31.65.40 : 1100)
Elapsed Time (seconds): 1320.256470
Total Directories Compared: 657
Total Directories Missing: 0
Total Directories Remirrored: 0
Total Files Compared: 120978
Total Files Missing: 0
Total Files Different: 3
Total Files Encrypted: 0
Total Files Remirrored: 1
Total Bytes Skipped: 0
Total Bytes Compared: 18527203678
Total Bytes Missing: 0
Total Bytes Different: 17
Total Bytes Remirrored: 17
Related links and directory attributes have been adjusted.
----- END OF VERIFICATION -----
```

- **Timestamp**—The last modified date and time of the file
- **Size**—The size, in bytes, of the file
- **Mask**—The attributes associated with the file. See further details below.
- **Security descriptors**—The NTFS file permissions of the file. If the file permissions are different, the message "Security descriptors are different" will be logged. If the file permissions are the same, nothing will be logged.
- **Bytes out of sync**—The number of bytes that are not synchronized between the file on the source and the file on the target. If the data in the file is identical, the message "0 BYTES OUT OF SYNC" will be logged. If the file is different, the message will indicate how many bytes were different. This message does not indicate that the file was remirrored during the verify.

The mask must be converted in order to determine what attributes are assigned to a file. The mask is a hexadecimal number corresponding to a binary number that indicates what the attributes are. Using the following steps, you can determine how the mask corresponds to the attributes of a file.

1. Each mask begins with 0x. Identify the hexadecimal number after the constant 0x. For example, if the mask is 0x23, then the hexadecimal number you are interested in is 23. The hexadecimal number may be up to four digits.
2. Convert the hexadecimal number to its 16-digit binary equivalent. You can use the Windows calculator for this conversion.
 - a. Select **Calculator** from your **Accessories** program or apps group.
 - b. Switch to scientific view, if it is not already in that view, by selecting **View, Scientific**.
 - c. Select **Hex**.
 - d. Enter the hexadecimal number, for example 23, as specified in your verification log.
 - e. Select **Bin** and the hexadecimal number will change to the binary equivalent.
 - f. Pad the beginning of the binary equivalent with zeroes (0) so that the number is 16 digits long. For example, hexadecimal number 23 converts to 100011, so the 16-digit binary equivalent would be 0000000000100011.
3. Determine what number (0 or 1) appears in each position of the binary number. Because binary numbers count from right to left, start with position 1 on the right.
 - 1—Read only
 - 2—Hidden
 - 3—None
 - 4—System
 - 5—Directory
 - 6—Archive
 - 7—Encrypted
 - 8—Normal
 - 9—Temporary
 - 10—Sparse file
 - 11—Reparse point
 - 12—Compressed
 - 13—Offline
 - 14—Not content indexed
 - 15—None
 - 16—None
4. Using the list above, identify those attributes that are enabled by those positions equal to one (1). The positions equal to zero (0) are disabled and that attribute does not apply. So hexadecimal number 23, which converted to 0000000000100011, indicates read only, hidden, and archive. Another example might be mask 0x827 which converted to binary is 0000100000100111. Positions 1-3, 6, and 12 are all enabled which indicates the file is read only, hidden, archive, and compressed.



Files that were replicated with the **Replicate NTFS security attributes by name** feature enabled, will be identified as different in the log file because of the local name attribute. The files will be the same.

Viewing server events

Highlight a server on the **Servers** page and click **View Server Events** from the toolbar. The **View Server Events** page displays the same messages that are logged to the Windows Event Viewer. The list of events are displayed in the top pane of the page, although the description is limited. When you highlight an event, the event details, including the full description, are displayed in the bottom pane of the page.

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


You can filter the events displayed by using the **Filter** drop-down list or the **View Warning Events** and **View Error Events** toolbar buttons. To clear a filter, select **All events** in the **Filter** drop-down list. See the *Reference Guide* for a complete list of Windows event messages.

Viewing server logs

You can view the engine and Management Service logs using either of these two methods.

- On the **Servers** page, highlight a server in the list and click **View Server Logs** from the toolbar.
- On the **Jobs** page, right-click a job and select **View Logs**. Select either the source server log or the target server log.

Separate logging windows allow you to continue working in the Double-Take Console while monitoring log messages. You can open multiple logging windows for multiple servers. When the Double-Take Console is closed, all logging windows will automatically close.



The following table identifies the controls and the table columns in the **Server logs** window.

Start 

This button starts the addition and scrolling of new messages in the window.

Pause 

This button pauses the addition and scrolling of new messages in the window. This is only for the **Server logs** window. The messages are still logged to their respective files on the server.

Copy 

This button copies the messages selected in the **Server logs** window to the Windows clipboard.

Clear 

This button clears the **Server logs** window. The messages are not cleared from the respective files on the server. If you want to view all of the messages again, close and reopen the **Server logs** window.

Filter

From the drop-down list, you can select to view all log messages or only those messages from the Double-Take log or the Double-Take Management Service log.

Time

This column in the table indicates the date and time when the message was logged.

Description

This column in the table displays the actual message that was logged.

Service

This column in the table indicates if the message is from the Double-Take log or the Double-Take Management Service log.

Managing snapshots

Use the instructions below to manage the snapshots that Double-Take has taken.

1. From the **Jobs** page, highlight the job and click **Manage Snapshots** in the toolbar.
2. You will see the list of snapshots, if any, associated with the job.
 - **Scheduled**—This snapshot was taken as part of a periodic snapshot.
 - **Deferred**—This snapshot was taken as part of a periodic snapshot, although it did not occur at the specified interval because the job between the source and target was not in a good state.
 - **Manual**—This snapshot was taken manually by a user.
3. Click **Take Snapshot** to create a new snapshot for the job.
4. If there is a snapshot that you no longer need, highlight it in the list and click **Delete**.
5. When you have completed your snapshot management, click **Close**.



If you have already failed over, the failover process will remove any Double-Take snapshots from the list. You will need to manage them manually using VSS. See your VSS documentation for more details.

Snapshot states

For some job types, when Double-Take transitions from a good state to a bad state, it will automatically attempt to take a snapshot of the data before it leaves the good state and enters the bad state. For example, if your data is in a good state and you start a mirror, before the mirror is started, Double-Take will automatically take a snapshot of the target. In the event the mirror fails to complete, you will have a snapshot of the data on the target when it was in its last good state. Only one automatic snapshot per job is maintained on the target. When an automatic snapshot is taken, it replaces any previous automatic snapshots.

A snapshot may not necessarily be useful if the data on the target is in a bad state. You only want snapshots of data that is in a good state. Therefore, you need to understand when the data is in a good or bad state.

- **Mirror started**
 - **State**—Bad
 - **Description**—Mirroring has started, but is not complete. The data on the source and target will not be synchronized until the mirror is complete.
 - **Automatic action taken for scheduled and automatic snapshots**—Scheduled and automatic snapshots will be delayed until the mirror is complete before taking a snapshot.
 - **User interaction required for manual snapshots**—Wait until the mirror is complete and the data is in a good state, then take a manual snapshot.
- **Mirror stopped**
 - **State**—Bad
 - **Description**—Mirroring has stopped without completing. The data on the source and target will not be synchronized until the mirror is complete.
 - **Automatic action taken for scheduled and automatic snapshots**—Scheduled and automatic snapshots will be delayed until the mirror has been restarted and is complete before taking a snapshot.
 - **User interaction required for manual snapshots**—Restart the mirror, wait until it is complete and the data is in a good state, and then take a manual snapshot.
- **Mirror complete**
 - **State**—Good
 - **Description**—Because the mirror is complete, the data on the source and target is synchronized. Double-Take will take a snapshot while the data is in a good state.
 - **Automatic action taken for scheduled and automatic snapshots**—Scheduled and automatic snapshots will occur normally.
 - **User interaction required for manual snapshots**—Manual snapshots can be taken normally.
- **Write operation retried**
 - **State**—Good
 - **Description**—An operation cannot be written to the hard drive on the target. For example, the file could be in use by another application on the target.
 - **Automatic action taken for scheduled and automatic snapshots**—Scheduled and automatic snapshots will occur normally, although the operation that is being retried will not be included in the snapshot.

- **User interaction required for manual snapshots**—Manual snapshots can be taken normally, although the operation that is being retried will not be included in the snapshot.
- **Write operation dropped**
 - **State**—Bad
 - **Description**—An operation could not be written to the hard drive on the target, even after multiple retries. For example, the file could be in use by another application on the target.
 - **Automatic action taken for scheduled and automatic snapshots**—An automatic snapshot will be taken just prior to the operation being dropped. Scheduled snapshots will be delayed until the target data is back in a good state.
 - **User interaction required for manual snapshots**—Start a mirror, wait until it is complete and the data is in a good state, and then take a manual snapshot.
- **Write operation succeeded**
 - **State**—Good
 - **Description**—An operation that was retrying on the target has been successfully written to the hard drive.
 - **Automatic action taken for scheduled and automatic snapshots**—Scheduled and automatic snapshots will occur normally.
 - **User interaction required for manual snapshots**—Manual snapshots can be taken normally.
- **Target restarted with job persistence**
 - **State**—Good
 - **Description**—The target service was able to persist job information prior to restarting.
 - **Automatic action taken for scheduled and automatic snapshots**—Scheduled and automatic snapshots will occur normally.
 - **User interaction required for manual snapshots**—Manual snapshots can be taken normally.
- **Target restarted without job persistence**
 - **State**—Bad
 - **Description**—The target service has been restarted and was unable to persist job information, therefore, operations that were in the queue have been lost.
 - **Automatic action taken for scheduled and automatic snapshots**—An automatic snapshot will be taken after the target restarts, if the target data was in a good state prior to the target restart and the job is configured to auto-remirror on auto-reconnect. Scheduled snapshots will be delayed until the target data is back in a good state.
 - **User interaction required for manual snapshots**—Start a mirror, wait until it is complete and the data is in a good state, and then take a manual snapshot.
- **Restore required**
 - **State**—Good or bad
 - **Description**—The data on the target no longer matches the data on the source because of a failover. This does not necessarily mean that the data on the target is bad.
 - **Automatic action taken for scheduled and automatic snapshots**—Scheduled and automatic snapshots will be delayed until a restore is completed or the restore required state is overruled by a mirror. Once the restoration or mirror is complete, automatic and scheduled snapshots will occur normally.

- **User interaction required for manual snapshots**—Restore the target data back to the source or override the restore required state by performing a mirror. Once the restoration or mirror is complete, manual snapshots can be taken normally.
- **Snapshot reverted**
 - **State**—Good or bad
 - **Description**—The data on the target no longer matches the data on the source because a snapshot has been applied on the target. This does not necessarily mean that the data on the target is bad.
 - **Automatic action taken for scheduled and automatic snapshots**—Scheduled and automatic snapshots will be delayed until a restore is completed or the snapshot reverted state is overruled by a mirror. Once the restoration or mirror is complete, automatic and scheduled snapshots will occur normally.
 - **User interaction required for manual snapshots**—Restore the target data back to the source or override the snapshot reverted state by performing a mirror. Once the restoration or mirror is complete, manual snapshots can be taken normally.
- **Restore complete**
 - **State**—Good
 - **Description**—Because the restoration is complete, the data on the source and target is synchronized.
 - **Automatic action taken for scheduled and automatic snapshots**—Scheduled and automatic snapshots will occur normally.
 - **User interaction required for manual snapshots**—Manual snapshots can be taken normally.

To be completely assured that your data on the target is good, automatic and scheduled snapshots only occur when the data is in a good Double-Take state. However, manual snapshots can be taken during any state. There are instances when you may want to take a manual snapshot, even if the target data is in a bad state. For example, if you drop an operation, that does not necessarily mean your data on the target is corrupt or the target would be unable to stand in for the source in the event of a failure. A snapshot of a bad state may be useful and usable, depending on your environment. If your source is a file server and an operation has been dropped, it is just one user file that is out-of-date. All of the remaining target files are intact and can be accessed in the event of a failure. However, if your source is an application server and an operation has been dropped, that one file could cause the application not to start on the target in the event of a failure. In these cases, manual snapshots of a bad state depend on the context of your environment.

Firewalls

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

- **Double-Take ports**—Ports 6320, 6325, and 6326 are used for Double-Take communications and must be open on your firewall. Open UDP and TCP for both inbound and outbound traffic.
- **Microsoft WMI and RPC ports**—Some features of Double-Take and the Double-Take Console 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.
- **Microsoft File Share and Directory ports**—Double-Take push installations will also rely on File Share and Directory ports, which must be open on your firewall. Check your Microsoft documentation if you need to modify these ports.
 - Microsoft File Share uses ports 135 through 139 for TCP and UDP communications.
 - Microsoft Directory uses port 445 for TCP and UDP communications.

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

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

Chapter 5 Protection

Create a DR protection job when you want to protect data or an entire server to a target repository server.

- *Creating a protection job* on page 72—This section includes step-by-step instructions for creating a protection job.
- *Managing and controlling protection jobs* on page 87—You can view overview status information about your protection job.
 - *Viewing protection job details* on page 98—You can view detailed status information and statistics for your protection job.
 - *Validating a protection job* on page 102—Over time, you may want to confirm that any changes in your network or environment have not impacted your job. Use these instructions to validate an existing job.
 - *Editing a protection job* on page 103—You can view and edit the settings for an existing job.
 - *Viewing a protection job log* on page 105—You can view the log associated with a specific job.


Creating a protection job

Use the following instructions to create a data protection or full server protection job.

1. From the **Servers** page, right-click the server you want to protect and select **Protect**. You can also highlight a server, click **Create a New Job** in the toolbar, then select **Protect**.
2. Choose the type of workload that you want to protect. Under **Server Workloads**, in the **Workload types** pane, select **Full Server DR Protection** if you want to protect the entire source server, or select **Files and Folders DR Protection** if you want to protect specific data on the source server. In the **Workload items** pane, select the volumes on the source that you want to protect.

If the workload you are looking for is not displayed, enable **Show all workload types**. The workload types in gray text are not available for the source server you have selected. Hover your mouse over an unavailable workload type to see a reason why this workload type is unavailable for the selected source.

3. To select your files and folders in more detail, click the **Replication Rules** heading and expand the volumes under **Folders**.



Volumes and folders with a green highlight are included completely. Volumes and folders highlighted in light yellow are included partially, with individual files or folders included. If there is no highlight, no part of the volume or folder is included. To modify the items selected, highlight a volume, folder, or file and click **Add Rule**. Specify if you want to **Include** or **Exclude** the item. Also, specify if you want the rule to be recursive, which indicates the rule should automatically be applied to the subdirectories of the specified path. If you do not select **Recursive**, the rule will not be applied to subdirectories.

You can also enter wildcard rules, however you should do so carefully. Rules are applied to files that are closest in the directory tree to them. If you have rules that include multiple folders, an exclusion rule with a wild card will need to be added for each folder that it needs applied to. For example, if you want to exclude all .log files from D:\ and your rules include D:\, D:\Dir1 , and

D:\Dir2, you would need to add the exclusion rule for the root and each subfolder rule. So you will need to add exclude rules for D:*.log , D:\Dir1*.log, and D:\Dir2*.log.

If you need to remove a rule, highlight it in the list at the bottom and click **Remove Rule**. Be careful when removing rules. Double-Take may create multiple rules when you are adding directories. For example, if you add E:\Data to be included in protection, then E:\ will be excluded. If you remove the E:\ exclusion rule, then the E:\Data rule will be removed also.



If you return to this page using the **Back** button in the job creation workflow, your **Workload Types** selection will be rebuilt, potentially overwriting any manual replication rules that you specified. If you do return to this page, confirm your **Workload Types** and **Replication Rules** are set to your desired settings before proceeding forward again.

4. Click **Next** to continue.
5. Choose your target repository server. This is the server that will store the image of the source or the image of the protected data from the source.



- **Current Servers**—This list contains the servers currently available in your console session. Servers that are not licensed for the workflow you have selected and those not applicable to the workload type you have selected will be filtered out of the list. Select your target server from the list. If the server you are looking for is not displayed, enable **Show all servers**. The servers in red are not available for the source server or workload type you have selected. Hover your mouse over an unavailable server to see a reason why this server is unavailable.
- **Find a New Server**—If the server you need is not in the **Current Servers** list, click the **Find a New Server** heading. From here, you can specify a server along with credentials for logging in to the server. If necessary, you can click **Browse** to select a server from a network drill-down list.



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

When specifying credentials for a new server, specify a user that is a member of the local Double-Take Admin and local administrator security groups.

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



You may be prompted for a route from the target to the source. This route, and a port if you are using a non-default port, is used so the target can communicate with the source to build job options. This dialog box will be displayed, only if needed.

7. You have many options available for your DR protection job. Configure those options that are applicable to your environment.

Go to each page identified below to see the options available for that section of the **Set Options** page. After you have configured your options, continue with the next step on page 86.

- *General* on page 75
- *Image Protection* on page 76
- *Mirror, Verify & Orphaned Files* on page 77
- *Network Route* on page 80
- *Snapshots* on page 81
- *Compression* on page 82
- *Bandwidth* on page 83
- *Scripts* on page 85

General



General

Job name:
alpha to beta

For the **Job name**, specify a unique name for your job.

Image Protection

Volume	Virtual Disk	Virtual Disk Path	Disk Name
C	Create new disk	C:\Server Images\Alpha	Alpha_C.vhdx

- **Image name**—Specify the name of the image that will store the protected data from the source.
- **Virtual Disk**—For each volume you are protecting, a virtual disk will be created on the target that contains the protected data from the source. Specify if you want to create a new disk or use an existing disk.

Reusing a virtual disk can be useful for pre-staging data on a LAN and then relocating the virtual disk to a remote site after the initial mirror is complete. You save time by skipping the virtual disk creation steps and performing a difference mirror instead of a full mirror. With pre-staging, less data will need to be sent across the wire initially. In order to use an existing virtual disk, it must be a valid virtual disk and it cannot be attached to a virtual machine. Additionally, the disk cannot be a differencing disk.

In a WAN environment, you may want to take advantage of using an existing disk by using a process similar to the following.

- a. Create a job in a LAN environment, letting Double-Take create the virtual disk for you.
 - b. Complete the mirror process locally.
 - c. Delete the job and when prompted, do not delete the virtual disks.
 - d. Copy or move the .vhd virtual files to your remote site.
 - e. Create a new protection job for the same source and reuse your existing disk.
- **Virtual Disk Path**—Specify the location on the target where you want to store the virtual disk that contains the protected data from the source.
 - **Disk Name**—Specify the name of the virtual disk file.

Mirror, Verify & Orphaned Files

Mirror, Verify & Orphaned Files

Mirror Options

Choose a comparison method and whether to mirror the entire file or only the bytes that differ in each file.

Compare file attributes and data. Send the attributes and bytes that differ. (Recommended for databases)

Verification Options

Enable scheduled verification

Verify on this interval: 1 Days

Begin immediately

Begin at this time: 9/26/2016 10:21:53 AM

Report and comparison options

Report only

Report and mirror files

Compare file attributes and data

General Options

Calculate size of protected data upon connection

Delete orphaned files

- **Mirror Options**—Choose a comparison method and whether to mirror the entire file or only the bytes that differ in each file.



If you are using a data protection job and are using a database application or are protecting a domain controller, do not use the compare file attributes only options unless you know for certain that you need it. With database applications and because domain controllers store their data in a database, it is critical that all files, not just some of the files, are mirrored. In this case, you should compare both the attributes and the data.

- **Do not compare files. Send the entire file.**—Double-Take will not perform any comparisons between the files on the source and target. All files will be mirrored to the target, sending the entire file. This is equivalent to selecting the mirror all files option prior to Double-Take version 7.1.
- **Compare file attributes. Send the entire file.**—Double-Take will compare file attributes and will mirror those files that have different attributes, sending the entire file. This is similar to selecting the mirror different files and the only if source is newer options prior to Double-Take version 7.1. This option is not available for full server protection jobs.
- **Compare file attributes. Send the attributes and bytes that differ.**—Double-Take will compare file attributes and will mirror only the attributes and bytes that are different. This is equivalent to selecting the mirror different files, only if source is newer, and use block checksum options prior to Double-Take version 7.1. This option is not available for full server protection jobs.
- **Compare file attributes and data. Send the attributes and bytes that differ.**—Double-Take will compare file attributes and the file data and will mirror only the attributes and bytes that are different. This is equivalent to selecting the mirror different files and use block checksum options prior to Double-Take version 7.1. If you are using a database application on your source, select this option.

- **Verification Options**—Choose if you want to periodically confirm that the source replica data on the target is identical to the actual data on the source. Verification creates a log file detailing what was verified as well as which files are not synchronized. If the data is not the same, you can automatically initiate a remirror, if configured. The remirror ensures data integrity between the source and target.
-



Because of the way the Windows Cache Manager handles memory, machines that are doing minimal or light processing may have file operations that remain in the cache until additional operations flush them out. This may make Double-Take files on the target appear as if they are not synchronized. When the Windows Cache Manager releases the operations in the cache on the source and target, the files will be updated on the target.

- **Enable scheduled verification**—When this option is enabled, Double-Take will verify the source replica data on the target.
 - **Verify on this interval**—Specify the interval between verification processes.
 - **Begin immediately**—Select this option if you want to start the verification schedule immediately after the job is established.
 - **Begin at this time**—Select this option if you want to start the verification schedule at the specified date and time.
 - **Report only**—Select this option if you only want to generate a verification report. With this option, no data that is found to be different will be mirrored to the target. Choose how you want the verification to compare the files.
 - **Report and mirror files**—Select this option if you want to generate a verification report and mirror data that is different to the target. Select the comparison method and type of mirroring you want to use. See the previous mirroring methods described under *Mirror Options*.
-



If you are using SQL to create snapshots of a SQL database, the verification report will report the file size of the snapshot files on the source and target as different. This is a reporting issue only. The snapshot file is mirrored and replicated completely to the target.

If you are using HP StorageWorks File Migration Agent, migrated files will incorrectly report modified time stamp differences in the verification report. This is a reporting issue only.

- **General Options**—Choose your general mirroring options. The calculate size option is only available for data protection and recovery. It is not available for full server protection and recovery.
 - **Calculate size of protected data upon connection**—Specify if you want Double-Take to determine the mirroring percentage calculation based on the amount of data being protected. If you enable this option, the calculation will begin when mirroring begins. For the initial mirror, the percentage will display after the calculation is complete, adjusting to the amount of the mirror that has completed during the time it

took to complete the calculation. Subsequent mirrors will initially use the last calculated size and display an approximate percentage. Once the calculation is complete, the percentage will automatically adjust down or up to indicate the amount that has been completed. Disabling calculation will result in the mirror status not showing the percentage complete or the number of bytes remaining to be mirrored.



The calculated amount of protected data may be slightly off if your data set contains compressed or sparse files.

This option is not available for full server protection jobs.

- **Delete orphaned files**—An orphaned file is a file that exists in the replica data on the target, but does not exist in the protected data on the source. This option specifies if orphaned files should be deleted on the target.
-



Orphaned file configuration is a per target configuration. All jobs to the same target will have the same orphaned file configuration.


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

The orphaned file feature does not delete alternate data streams. To do this, use a full mirror, which will delete the additional streams when the file is re-created.

If you want to move orphaned files rather than delete them, you can configure this option along with the move deleted files feature to move your orphaned files to the specified deleted files directory. See *Target server properties* on page 54 for more information.

During a mirror, orphaned file processing success messages will be logged to a separate orphaned file log on the source. This keeps the Double-Take log from being overrun with orphaned file success processing messages. Orphaned files processing statistics and any errors in orphaned file processing will still be logged to the Double-Take log, and during difference mirrors, verifications, and restorations, all orphaned file processing messages are logged to the Double-Take log. The orphaned file log is located in the **Logging folder** specified for the source. See *Log file properties* on page 59 for details on the location of that folder. The orphaned log file is appended to during each orphaned file processing during a mirror, and the log file will be a maximum of 50 MB.

Network Route



Network Route

Send data to the target server using this route:

10.10.10.30

By default, Double-Take will select a target route for transmissions. If desired, specify an alternate route on the target that the data will be transmitted through. This allows you to select a different route for Double-Take traffic. For example, you can separate regular network traffic and Double-Take traffic on a machine with multiple IP addresses. You can also select or manually enter a public IP address (which is the public IP address of the server's router) if you are using a NAT environment. If you enter a public IP addresses, you will see additional fields allowing you to disable the default communication ports and specify other port numbers to use, allowing the target to communicate through a router. The **Management Service port** is used to persist the source share configuration when shares are being protected. The **Replication Service port** is used for data transmission.



If you change the IP address on the target which is used for the target route, you will be unable to edit the job. If you need to make any modifications to the job, it will have to be deleted and re-created.

Snapshots

The screenshot shows a configuration window titled "Snapshots". It has a header bar with a back arrow and the title. Below the header, there are three main options:

- Enable scheduled snapshots
- Take snapshots on this interval: 1 Hours (selected)
- Begin immediately
- Begin at this time: 1/11/2016 3:45:38 PM (selected)

A snapshot is an image of the image of the source or of the image of the protected data from the source. In other words, your DR job will create an image of the source or an image of the protected data from the source. If you take a snapshot of that image, you are creating an image of an image. A snapshot is an image taken at a single point in time. You can recover to a snapshot. However, you cannot access the snapshot in VSS to recover any data.

Turn on **Enable scheduled snapshots** if you want Double-Take to take snapshots automatically at set intervals.

- **Take snapshots on this interval**—Specify the interval (in days, hours, or minutes) for taking snapshots.
- **Begin immediately**—Select this option if you want to start taking snapshots immediately after the protection job is established.
- **Begin at this time**—Select this option if you want to start taking snapshots at a later date and time. Specify the date and time parameters to indicate when you want to start.



See *Managing snapshots* on page 66 for details on taking manual snapshots and deleting snapshots.

Snapshots are stored inside the image of the source on the target, so be sure that you have selected a volume that has enough space for the image of the source with its snapshots, as well as any other source images and their snapshots.

You may want to set the size limit on how much space snapshots can use. See your VSS documentation for more details.

Compression



To help reduce the amount of bandwidth needed to transmit Double-Take data, compression allows you to compress data prior to transmitting it across the network. In a WAN environment this provides optimal use of your network resources. If compression is enabled, the data is compressed before it is transmitted from the source. When the target receives the compressed data, it decompresses it and then writes it to disk. You can set the level from **Minimum** to **Maximum** to suit your needs.


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

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



All jobs from a single source connected to the same IP address on a target will share the same compression configuration.

Bandwidth



Bandwidth

Do not limit bandwidth

Use a fixed limit (bytes per second):

Preset bandwidth: Bandwidth (bytes per second):

Use scheduled limits:

Days	Start	End	Limit
Every day	6:00 AM	6:00 PM	1250000
Every day	6:00 PM	6:00 AM	5625000

 Bandwidth is not limited during unscheduled times.

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



All jobs from a single source connected to the same IP address on a target will share the same bandwidth configuration.

- **Do not limit bandwidth**—Double-Take will transmit data using 100% bandwidth availability.
- **Use a fixed limit**—Double-Take will transmit data using a limited, fixed bandwidth. Select a **Preset bandwidth** limit rate from the common bandwidth limit values. The **Bandwidth** field will automatically update to the bytes per second value for your selected bandwidth. This is the maximum amount of data that will be transmitted per second. If desired, modify the bandwidth using a bytes per second value. The minimum limit should be 3500 bytes per second.
- **Use scheduled limits**—Double-Take will transmit data using a dynamic bandwidth based on the schedule you configure. Bandwidth will not be limited during unscheduled times.
 - **New**—Click **New** to create a new scheduled bandwidth limit. Specify the following information.
 - **Daytime entry**—Select this option if the start and end times of the bandwidth window occur in the same day (between 12:01 AM and midnight). The start time must occur before the end time.
 - **Overnight entry**—Select this option if the bandwidth window begins on one day and continues past midnight into the next day. The start time must be later than the end time, for example 6 PM to 6 AM.
 - **Day**—Enter the day on which the bandwidth limiting should occur. You can pick a specific day of the week, **Weekdays** to have the limiting occur Monday through Friday, **Weekends** to have the limiting occur Saturday and Sunday, or **Every day** to have the limiting repeat on all days of the week.
 - **Start time**—Enter the time to begin bandwidth limiting.

- **End time**—Enter the time to end bandwidth limiting.
 - **Preset bandwidth**—Select a bandwidth limit rate from the common bandwidth limit values. The **Bandwidth** field will automatically update to the bytes per second value for your select bandwidth.
 - **Bandwidth**—If desired, modify the bandwidth using a bytes per second value. The minimum limit should be 3500 bytes per second.
 - **Edit**—Click **Edit** to modify an existing scheduled bandwidth limit.
 - **Delete**—Click **Delete** to remove a scheduled bandwidth limit.
-



If you change your job option from **Use scheduled limits** to **Do not limit bandwidth** or **Use a fixed limit**, any schedule that you created will be preserved. That schedule will be reused if you change your job option back to **Use scheduled limits**.

You can manually override a schedule after a job is established by selecting **Other Job Options, Set Bandwidth**. If you select **No bandwidth limit** or **Fixed bandwidth limit**, that manual override will be used until you go back to your schedule by selecting **Other Job Options, Set Bandwidth, Scheduled bandwidth limit**. For example, if your job is configured to use a daytime limit, you would be limited during the day, but not at night. But if you override that, your override setting will continue both day and night, until you go back to your schedule. See the *Managing and controlling jobs* section for your job type for more information on the **Other Job Options**.

Scripts

The screenshot shows a window titled "Scripts" with three sections:

- Mirror Start**: Script file: `c:\scripts\mirrorstart.bat`; Arguments: (empty); Allow script to interact with desktop; Delay until script completes.
- Mirror Complete**: Script file: (empty); Arguments: (empty); Allow script to interact with desktop; Delay until script completes.
- Mirror Stop**: Script file: `c:\scripts\mirrorcomplete.bat`; Arguments: `arg1`; Allow script to interact with desktop; Delay until script completes.

You can customize mirroring by running scripts on the target at pre-defined points in the mirroring process. Scripts may contain any valid Windows command, executable, or batch file. The scripts are processed using the same account running the Double-Take service, unless you have identified a specific account through the server's properties. See *Script credentials* on page 58. There are three types of mirroring scripts.

- **Mirror Start**—This script starts when the target receives the first mirror operation. In the case of a difference mirror, this may be a long time after the mirror is started because the script does not start until the first different data is received on the target. If the data is synchronized and a difference mirror finds nothing to mirror, the script will not be executed. Specify the full path and name of the **Script file**.
- **Mirror Complete**—This script starts when a mirror is completed. Because the mirror statistics may indicate a mirror is at 99-100% when it is actually still processing (for example, if files were added after the job size was calculated, if there are alternate data streams, and so on), the script will not start until all of the mirror data has been completely processed on the target. Specify the full path and name of the **Script file**.
- **Mirror Stop**—This script starts when a mirror is stopped, which may be caused by an auto-disconnect occurring while a mirror is running, the service is shutdown while a mirror is running, or if you stop a mirror manually. Specify the full path and name of the **Script file**.
- **Arguments**—Specify a comma-separated list of valid arguments required to execute the script.
- **Allow script to interact with desktop**—Enable this option if you want the script processing to be displayed on the screen. Otherwise, the script will execute silently in the background. If you enable this option, the account running the script must have permission to interact with the desktop.
- **Delay until script completes**—Enable this option if you want to delay the mirroring process until the associated script has completed. If you select this option, make sure your script handles errors, otherwise the mirroring process may never complete if the process is waiting on a script that cannot complete.
- **Test**—You can test your script manually by clicking **Test**. Your script will be executed if you test it. If necessary, manually undo any changes that you do not want on your target after testing the script.



If you establish mirroring scripts for one job and then establish additional jobs to the same target using the same target path mapping, the mirroring scripts will automatically be applied to those subsequent jobs. If you select a different target path mapping, the mirroring scripts will have to be reconfigured for the new job(s).

8. Click **Next** to continue.
9. Double-Take validates that your source and target are compatible. The **Summary** page displays your options and validation items.

Errors are designated by a white X inside a red circle. Warnings are designated by a black exclamation point (!) inside a yellow triangle. A successful validation is designated by a white checkmark inside a green circle. You can sort the list by the icon to see errors, warnings, or successful validations together. Click on any of the validation items to see details. You must correct any errors before you can continue. Depending on the error, you may be able to click **Fix** or **Fix All** and let Double-Take correct the problem for you. For those errors that Double-Take cannot correct automatically, you will need to modify the source or target to correct the error, or you can select a different target. You must revalidate the selected servers, by clicking **Recheck**, until the validation check passes without errors.

If you receive a path transformation error during job validation indicating a volume does not exist on the target server, even though there is no corresponding data being protected on the source, you will need to manually modify your replication rules. Go back to the **Choose Data** page and under the **Replication Rules**, locate the volume from the error message. Remove any rules associated with that volume. Complete the rest of the workflow and the validation should pass.

Before a job is created, the results of the validation checks are logged to the Double-Take Management Service log on the target. After a job is created, the results of the validation checks are logged to the job log. See the Double-Take *Reference Guide* for details on the various Double-Take log files.

10. Once your servers have passed validation and you are ready to establish protection, click **Finish**, and you will automatically be taken to the **Jobs** page.



Jobs in a NAT environment may take longer to start.

Managing and controlling protection jobs

Click **Jobs** from the main Double-Take Console toolbar. The **Jobs** page allows you to view status information about your jobs. You can also control your jobs from this page.

The jobs displayed in the right pane depend on the server group folder selected in the left pane. Every job for each server in your console session is displayed when the **Jobs on All Servers** group is selected. If you have created and populated server groups (see *Managing servers* on page 24), then only the jobs associated with the server or target servers in that server group will be displayed in the right pane.

- *Overview job information displayed in the top right pane* on page 87
- *Detailed job information displayed in the bottom right pane* on page 90
- *Job controls* on page 92

Overview job information displayed in the top right pane

The top pane displays high-level overview information about your jobs. You can sort the data within a column in ascending and descending order. You can also move the columns to the left or right of each other to create your desired column order. The list below shows the columns in their default left to right order.

If you are using server groups, you can filter the jobs displayed in the top right pane by expanding the **Server Groups** heading and selecting a server group.

Column 1 (Blank)

The first blank column indicates the state of the job.

 A green circle with a white checkmark indicates the job is in a healthy state. No action is required.

 A yellow triangle with a black exclamation point indicates the job is in a pending or warning state. This icon is also displayed on any server groups that you have created that contain a job in a pending or warning state. Double-Take is working or waiting on a pending process or attempting to resolve the warning state.

 A red circle with a white X indicates the job is in an error state. This icon is also displayed on any server groups that you have created that contain a job in an error state. You will need to investigate and resolve the error.

 The job is in an unknown state.

Job

The name of the job

Source Server

The name of the source. This could be the name or IP address of your source.

Target Server

The name of the target. This could be the name or IP address of your target. For protection jobs, you will see (Repository) after the server name or IP address. This will help you distinguish this server and its job role.

Job Type

Each job type has a unique job type name. This job is either a Full Server to Windows DR Protection job or a Data Only for Windows DR Protection job, depending on if you are protecting the entire server or just specific data from the server. For a complete list of all job type names, press F1 to view the Double-Take Console online help.

Activity

There are many different **Activity** messages that keep you informed of the job activity. Most of the activity messages are informational and do not require any administrator interaction. If you see error messages, check the job details. Keep in mind that **Idle** indicates console to server activity is idle, not that your servers are idle.

Mirror Status

- **Calculating**—The amount of data to be mirrored is being calculated.
- **In Progress**—Data is currently being mirrored.
- **Waiting**—Mirroring is complete, but data is still being written to the target.
- **Idle**—Data is not being mirrored.
- **Paused**—Mirroring has been paused.
- **Stopped**—Mirroring has been stopped.
- **Removing Orphans**—Orphan files on the target are being removed or deleted depending on the configuration.
- **Verifying**—Data is being verified between the source and target.
- **Unknown**—The console cannot determine the status.

Replication Status

- **Replicating**—Data is being replicated to the target.
- **Ready**—There is no data to replicate.
- **Pending**—Replication is pending.
- **Stopped**—Replication has been stopped.
- **Out of Memory**—Replication memory has been exhausted.
- **Failed**—The Double-Take service is not receiving replication operations from the Double-Take driver. Check the Event Viewer for driver related issues.
- **Unknown**—The console cannot determine the status.

Transmit Mode

- **Active**—Data is being transmitted to the target.
- **Paused**—Data transmission has been paused.
- **Scheduled**—Data transmission is waiting on schedule criteria.
- **Stopped**—Data is not being transmitted to the target.

- **Error**—There is a transmission error.
- **Unknown**—The console cannot determine the status.

Operating System

The job type operating system

Detailed job information displayed in the bottom right pane

The details displayed in the bottom pane provide additional information for the job highlighted in the top pane. You can expand or collapse the bottom pane by clicking on the **Job Highlights** heading.

Name

The name of the job

Target data state

- **OK**—The data on the target is in a good state.
- **Mirroring**—The target is in the middle of a mirror process. The data will not be in a good state until the mirror is complete.
- **Mirror Required**—The data on the target is not in a good state because a remirror is required. This may be caused by an incomplete or stopped mirror or an operation may have been dropped on the target.
- **Busy**—The source is low on memory causing a delay in getting the state of the data on the target.
- **Not Loaded**—Double-Take target functionality is not loaded on the target server. This may be caused by a license key error.
- **Unknown**—The console cannot determine the status.

Mirror remaining

The total number of mirror bytes that are remaining to be sent from the source to the target. This value may be zero if you have enabled **Mirror only changed files when source reboots** on the *Server setup properties* on page 44.

Mirror skipped

The total number of bytes that have been skipped when performing a difference. These bytes are skipped because the data is not different on the source and target. This value may be zero if you have enabled **Mirror only changed files when source reboots** on the *Server setup properties* on page 44.

Replication queue

The total number of replication bytes in the source queue

Disk queue

The amount of disk space being used to queue data on the source

Recovery point latency

The length of time replication is behind on the target compared to the source. This is the time period of data that would be lost if a failure were to occur at the current time.

Bytes sent

The total number of mirror and replication bytes that have been transmitted to the target

Bytes sent (compressed)

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

Connected since

The date and time indicating when the current job was started.

Recent activity

Displays the most recent activity for the selected job, along with an icon indicating the success or failure of the last initiated activity. Click the link to see a list of recent activities for the selected job. You can highlight an activity in the list to display additional details about the activity.

Image name

The name of the image on the repository server

Additional information

Depending on the current state of your job, you may see additional information displayed to keep you informed about the progress and status of your job. If there is no additional information, you will see (None) displayed.

Job controls

You can control your job through the toolbar buttons available on the **Jobs** page. If you select multiple jobs, some of the controls will apply only to the first selected job, while others will apply to all of the selected jobs. For example, **View Job Details** will only show details for the first selected job, while **Stop** will stop protection for all of the selected jobs.

If you want to control just one job, you can also right click on that job and access the controls from the pop-up menu.

View Job Details

This button leaves the **Jobs** page and opens the **View Job Details** page.

Edit Job Properties

This button leaves the **Jobs** page and opens the **EditJob Properties** page.

Delete

Stops (if running) and deletes the selected jobs.

You have the option of deleting the images and virtual disks associated with your DR protection job.

- **Delete the associated images**—Select this option if you want to delete the job's recovery images. This option will keep the job's virtual disk, which can be reused later with another job, if desired. Keep in mind, if you delete the image, you will not be able to recover it. If you want to be able to recover to this image, do not select this option. Also, keep in mind that recovering to an image that no longer has an active job will only include data up to the point when you delete the job.
- **Delete the associated virtual disks**—Select this option if you want to delete the job's virtual disk. This option will delete the underlying .vhd file from your target repository server. Therefore, you will not be able to reuse it later with another job. If you want to be able to reuse the virtual disk, do not select this option.

Provide Credentials

Changes the login credentials that the job (which is on the target machine) uses to authenticate to the servers in the job. This button opens the Provide Credentials dialog box where you can specify the new account information and which servers you want to update. See *Providing server credentials* on page 37. You will remain on the **Jobs** page after updating the server credentials. If your servers use the same credentials, make sure you also update the credentials on the **Servers** page so that the Double-Take Console can authenticate to the servers in the console session. See *Managing servers* on page 24.

View Recent Activity

Displays the recent activity list for the selected job. Highlight an activity in the list to display additional details about the activity.

Start

Starts or resumes the selected jobs.

If you have previously stopped protection, the job will restart mirroring and replication.

If you have previously paused protection, the job will continue mirroring and replication from where it left off, as long as the Double-Take queue was not exhausted during the time the job was paused. If the Double-Take queue was exhausted during the time the job was paused, the job will restart mirroring and replication.

Also if you have previously paused protection, all jobs from the same source to the same IP address on the target will be resumed.

Pause

Pauses the selected jobs. Data will be queued on the source while the job is paused.

All jobs from the same source to the same IP address on the target will be paused.

Stop

Stops the selected jobs. The jobs remain available in the console, but there will be no mirroring or replication data transmitted from the source to the target. Mirroring and replication data will not be queued on the source while the job is stopped, requiring a remirror when the job is restarted. The type of remirror will depend on your job settings. Stopping a job will delete any Double-Take snapshots on the target.

Take Snapshot

Even if you have scheduled the snapshot process, you can run it manually any time. If an automatic or scheduled snapshot is currently in progress, Double-Take will wait until that one is finished before taking the manual snapshot.

Manage Snapshots

Allows you to manage your snapshots by taking and deleting snapshots for the selected job. See *Managing snapshots* on page 66 for more information.

Failover, Cutover, or Recover

Finishes the recovery process. Recovery does not apply to DR protection jobs.

Failback

Starts the failback process. Failback does not apply to DR jobs.

Restore

Starts the restoration process. Restore does not apply to DR jobs.

Reverse

Reverses protection. Reverse protection does not apply to DR jobs.

Recover Data

Recovers the selected DR job. You will be taken directly to the **Choose Source Image to Recover** page, which bypasses the first three pages of the recovery workflow (**Get Started**, **Double-Take DR**, and **Choose Source Repository Server**), automatically using the repository server of the selected protection job. See *Creating a recovery job* on page 108 for details on recovering a DR job.

Undo Failover or Cutover

Cancels a failover or cutover by undoing it. Undo failover does not apply to DR jobs.

View Job Log

Opens the job log. On the right-click menu, this option is called **View Logs**, and you have the option of opening the job log, source server log, or target server log.

Other Job Actions

Opens a small menu of other job actions. These job actions will be started immediately, but keep in mind that if you stop and restart your job, the job's configured settings will override any other job actions you may have initiated.

- **Mirroring**—You can start, stop, pause and resume mirroring for any job that is running.

When pausing a mirror, Double-Take stops queuing mirror data on the source but maintains a pointer to determine what information still needs to be mirrored to the target. Therefore, when resuming a paused mirror, the process continues where it left off.

When stopping a mirror, Double-Take stops queuing mirror data on the source and does not maintain a pointer to determine what information still needs to be mirrored to the target. Therefore, when starting a mirror that has been stopped, you will need to decide what type of mirror to perform.

- **Mirror Options**—Choose a comparison method and whether to mirror the entire file or only the bytes that differ in each file.



If you are using a data protection job and are using a database application or are protecting a domain controller, do not use the



compare file attributes only options unless you know for certain that you need it. With database applications and because domain controllers store their data in a database, it is critical that all files, not just some of the files, are mirrored. In this case, you should compare both the attributes and the data.

- **Do not compare files. Send the entire file.**—Double-Take will not perform any comparisons between the files on the source and target. All files will be mirrored to the target, sending the entire file. This is equivalent to selecting the mirror all files option prior to Double-Take version 7.1.
- **Compare file attributes. Send the entire file.**—Double-Take will compare file attributes and will mirror those files that have different attributes, sending the entire file. This is similar to selecting the mirror different files and the only if source is newer options prior to Double-Take version 7.1. This option is not available for full server protection jobs.
- **Compare file attributes. Send the attributes and bytes that differ.**—Double-Take will compare file attributes and will mirror only the attributes and bytes that are different. This is equivalent to selecting the mirror different files, only if source is newer, and use block checksum options prior to Double-Take version 7.1. This option is not available for full server protection jobs.
- **Compare file attributes and data. Send the attributes and bytes that differ.**—Double-Take will compare file attributes and the file data and will mirror only the attributes and bytes that are different. This is equivalent to selecting the mirror different files and use block checksum options prior to Double-Take version 7.1. If you are using a database application on your source, select this option.
- **Calculate size of protected data before mirroring**—Specify if you want Double-Take to determine the mirroring percentage calculation based on the amount of data being protected. If you enable this option, the calculation will begin when mirroring begins. For the initial mirror, the percentage will display after the calculation is complete, adjusting to the amount of the mirror that has completed during the time it took to complete the calculation. Subsequent mirrors will initially use the last calculated size and display an approximate percentage. Once the calculation is complete, the percentage will automatically adjust down or up to indicate the amount that has been completed. Disabling calculation will result in the mirror status not showing the percentage complete or the number of bytes remaining to be mirrored.



The calculated amount of protected data may be slightly off if your data set contains compressed or sparse files.

- **Verify**—Even if you have scheduled the verification process, you can run it manually any time a mirror is not in progress.

- **Report only**—Select this option if you only want to generate a verification report. With this option, no data that is found to be different will be mirrored to the target. Choose how you want the verification to compare the files.
- **Report and mirror files**—Select this option if you want to generate a verification report and mirror data that is different to the target. Select the comparison method and type of mirroring you want to use. See the previous mirroring methods described under *Mirror Options*.
- **Set Bandwidth**—You can manually override bandwidth limiting settings configured for your job at any time.
 - **No bandwidth limit**—Double-Take will transmit data using 100% bandwidth availability.
 - **Fixed bandwidth limit**—Double-Take will transmit data using a limited, fixed bandwidth. Select a **Preset bandwidth** limit rate from the common bandwidth limit values. The **Bandwidth** field will automatically update to the bytes per second value for your selected bandwidth. This is the maximum amount of data that will be transmitted per second. If desired, modify the bandwidth using a bytes per second value. The minimum limit should be 3500 bytes per second.
 - **Scheduled bandwidth limit**—If your job has a configured scheduled bandwidth limit, you can enable that schedule with this option.
- **Delete Orphans**—Even if you have enabled orphan file removal during your mirror and verification processes, you can manually remove them at any time.
- **Target**—You can pause the target, which queues any incoming Double-Take data from the source on the target. All active jobs to that target will complete the operations already in progress. Any new operations will be queued on the target until the target is resumed. The data will not be committed until the target is resumed. Pausing the target only pauses Double-Take processing, not the entire server.

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

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

Filter

Select a filter option from the drop-down list to only display certain jobs. You can display **Healthy jobs**, **Jobs with warnings**, or **Jobs with errors**. To clear the filter, select **All jobs**. If you have created and populated server groups, then the filter will only apply to the jobs associated with the server or target servers in that server group. See *Managing servers* on page 24.

Search

Allows you to search the source or target server name for items in the list that match the criteria you have entered.

Overflow Chevron 

Displays any toolbar buttons that are hidden from view when the window size is reduced.

Viewing protection job details

From the **Jobs** page, highlight the job and click **View Job Details** in the toolbar.

Review the following table to understand the detailed information about your job displayed on the **View Job Details** page.

Job name

The name of the job

Job type

Each job type has a unique job type name. This job is either a Full Server to Windows DR Protection job or a Data Only for Windows DR Protection job, depending on if you are protecting the entire server or just specific data from the server. For a complete list of all job type names, press F1 to view the Double-Take Console online help.

Health

-  The job is in a healthy state.
-  The job is in a warning state.
-  The job is in an error state.
-  The job is in an unknown state.

Activity

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

Connection ID

The incremental counter used to number connections. The number is incremented when a connection is created. The counter is reset if there are no existing jobs and the Double-Take service is restarted.

Transmit mode

- **Active**—Data is being transmitted to the target.
- **Paused**—Data transmission has been paused.
- **Scheduled**—Data transmission is waiting on schedule criteria.
- **Stopped**—Data is not being transmitted to the target.
- **Error**—There is a transmission error.
- **Unknown**—The console cannot determine the status.

Target data state

- **OK**—The data on the target is in a good state.
- **Mirroring**—The target is in the middle of a mirror process. The data will not be in a good state until the mirror is complete.
- **Mirror Required**—The data on the target is not in a good state because a remirror is required. This may be caused by an incomplete or stopped mirror or an operation may have been dropped on the target.
- **Busy**—The source is low on memory causing a delay in getting the state of the data on the target.
- **Not Loaded**—Double-Take target functionality is not loaded on the target server. This may be caused by a license key error.
- **Unknown**—The console cannot determine the status.

Target route

The IP address on the target used for Double-Take transmissions.

Compression

- **On / Level**—Data is compressed at the level specified.
- **Off**—Data is not compressed.

Encryption

- **On**—Data is being encrypted before it is sent from the source to the target.
- **Off**—Data is not being encrypted before it is sent from the source to the target.

Bandwidth limit

If bandwidth limiting has been set, this statistic identifies the limit. The keyword **Unlimited** means there is no bandwidth limit set for the job.

Connected since

The source server date and time indicating when the current job was started. This field is blank, indicating that a TCP/IP socket is not present, when the job is waiting on transmit options or if the transmission has been stopped. This field will maintain the date and time, indicating that a TCP/IP socket is present, when transmission has been paused.

Additional information

Depending on the current state of your job, you may see additional information displayed to keep you informed about the progress and status of your job. If there is no additional information, you will see (None) displayed.

Mirror status

- **Calculating**—The amount of data to be mirrored is being calculated.
- **In Progress**—Data is currently being mirrored.
- **Waiting**—Mirroring is complete, but data is still being written to the target.
- **Idle**—Data is not being mirrored.

- **Paused**—Mirroring has been paused.
- **Stopped**—Mirroring has been stopped.
- **Removing Orphans**—Orphan files on the target are being removed or deleted depending on the configuration.
- **Verifying**—Data is being verified between the source and target.
- **Unknown**—The console cannot determine the status.

Mirror percent complete

The percentage of the mirror that has been completed

Mirror remaining

The total number of mirror bytes that are remaining to be sent from the source to the target. This value may be zero if you have enabled **Mirror only changed files when source reboots** on the *Server setup properties* on page 44.

Mirror skipped

The total number of bytes that have been skipped when performing a difference. These bytes are skipped because the data is not different on the source and target. This value may be zero if you have enabled **Mirror only changed files when source reboots** on the *Server setup properties* on page 44.

Replication status

- **Replicating**—Data is being replicated to the target.
- **Ready**—There is no data to replicate.
- **Pending**—Replication is pending.
- **Stopped**—Replication has been stopped.
- **Out of Memory**—Replication memory has been exhausted.
- **Failed**—The Double-Take service is not receiving replication operations from the Double-Take driver. Check the Event Viewer for driver related issues.
- **Unknown**—The console cannot determine the status.

Replication queue

The total number of replication bytes in the source queue

Disk queue

The amount of disk space being used to queue data on the source

Bytes sent

The total number of mirror and replication bytes that have been transmitted to the target

Bytes sent compressed

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

Recovery point latency

The length of time replication is behind on the target compared to the source. This is the time period of data that would be lost if a failure were to occur at the current time.

Mirror start time

The UTC time when mirroring started

Mirror end time

The UTC time when mirroring ended

Total time for last mirror

The length of time it took to complete the last mirror process

Validating a protection job

Over time, you may want to confirm that any changes in your network or environment have not impacted your Double-Take job. Use these instructions to validate an existing job.

1. From the **Jobs** page, highlight the job and click **View Job Details** in the toolbar.
2. In the **Tasks** area on the right on the **View Job Details** page, click **Validate job properties**.
3. Double-Take validates that your source and target are compatible. The **Summary** page displays your options and validation items.

Errors are designated by a white X inside a red circle. Warnings are designated by a black exclamation point (!) inside a yellow triangle. A successful validation is designated by a white checkmark inside a green circle. You can sort the list by the icon to see errors, warnings, or successful validations together. Click on any of the validation items to see details. You must correct any errors before you can continue. Depending on the error, you may be able to click **Fix** or **Fix All** and let Double-Take correct the problem for you. For those errors that Double-Take cannot correct automatically, you will need to modify the source or target to correct the error, or you can select a different target. You must revalidate the selected servers, by clicking **Recheck**, until the validation check passes without errors.

Validation checks for an existing job are logged to the job log on the target server.

4. Once your servers have passed validation, click **Close**.

Editing a protection job

Use these instructions to edit a protection job.

1. From the **Jobs** page, highlight the job and click **View Job Details** in the toolbar.
2. In the **Tasks** area on the right on the **View Job Details** page, click **Edit job properties**. (You will not be able to edit a job if you have removed the source of that job from your Double-Take Console session or if you only have Double-Take monitor security access.)
3. You will see the same options available for your protection job as when you created the job, but you will not be able to edit all of them. If desired, edit those options that are configurable for an existing job. See *Creating a protection job* on page 72 for details on each job option.



Changing some options may require Double-Take to automatically disconnect, reconnect, and remirror the job.

If you have specified replication rules that exclude a volume at the root, that volume will be incorrectly added as an inclusion if you edit the job after it has been established. If you need to edit your job, modify the replication rules to make sure they include the proper inclusion and exclusion rules that you want.

4. If you want to modify the workload items or replication rules for the job, click **Edit workload or replication rules**.

In the **Server Workloads** section, you can add or remove any volumes from your protected data. If you add a new volume, you can create a new disk or use an existing disk after you click **OK** on the **Choose Data** page and are returned to the **Edit Job Properties** page. If you remove a volume, the image for the volume will be removed and you will not be able to recover that volume, including any snapshots for that volume. You can specify if you want to delete the virtual disk associated with the volume's image. If you delete the virtual disk, you will not be able to reuse that disk in a future DR protection job. If you want to be able to reuse the virtual disk, do not delete the virtual disk.

You can also modify the specific **Replication Rules** for your job. Volumes and folders with a green highlight are included completely. Volumes and folders highlighted in light yellow are included partially, with individual files or folders included. If there is no highlight, no part of the volume or folder is included. To modify the items selected, highlight a volume, folder, or file and click **Add Rule**. Specify if you want to **Include** or **Exclude** the item. Also, specify if you want the rule to be recursive, which indicates the rule should automatically be applied to the subdirectories of the specified path. If you do not select **Recursive**, the rule will not be applied to subdirectories. You can also enter wildcard rules, however you should do so carefully. Rules are applied to files that are closest in the directory tree to them. If you have rules that include multiple folders, an exclusion rule with a wild card will need to be added for each folder that it needs applied to. For example, if you want to exclude all .log files from D:\ and your rules include D:\, D:\Dir1, and D:\Dir2, you would need to add the exclusion rule for the root and each subfolder rule. So you will need to add exclude rules for D:*.log, D:\Dir1*.log, and D:\Dir2*.log. If you need to remove a rule, highlight it in the list at the bottom and click **Remove Rule**. Be careful when removing rules. Double-Take may create multiple rules when you are adding directories. For

example, if you add E:\Data to be included in protection, then E:\ will be excluded. If you remove the E:\ exclusion rule, then the E:\Data rule will be removed also.

Click **OK** to return to the **Edit Job Properties** page.

5. Click **Next** to continue.
6. Double-Take validates that your source and target are compatible. The **Summary** page displays your options and validation items.

Errors are designated by a white X inside a red circle. Warnings are designated by a black exclamation point (!) inside a yellow triangle. A successful validation is designated by a white checkmark inside a green circle. You can sort the list by the icon to see errors, warnings, or successful validations together. Click on any of the validation items to see details. You must correct any errors before you can continue. Depending on the error, you may be able to click **Fix** or **Fix All** and let Double-Take correct the problem for you. For those errors that Double-Take cannot correct automatically, you will need to modify the source or target to correct the error, or you can select a different target. You must revalidate the selected servers, by clicking **Recheck**, until the validation check passes without errors.


If you receive a path transformation error during job validation indicating a volume does not exist on the target server, even though there is no corresponding data being protected on the source, you will need to manually modify your replication rules. Go back to the **Choose Data** page and under the **Replication Rules**, locate the volume from the error message. Remove any rules associated with that volume. Complete the rest of the workflow and the validation should pass.

Before a job is created, the results of the validation checks are logged to the Double-Take Management Service log on the target. After a job is created, the results of the validation checks are logged to the job log. See the Double-Take *Reference Guide* for details on the various Double-Take log files.

7. Once your servers have passed validation and you are ready to update your job, click **Finish**.

Viewing a protection job log

You can view a job log file through the Double-Take Console by selecting **View Job Log** from the toolbar on the **Jobs** page. Separate logging windows allow you to continue working in the Double-Take Console while monitoring log messages. You can open multiple logging windows for multiple jobs. When the Double-Take Console is closed, all logging windows will automatically close.



Time	Description
9/27/2016 7:55:21 AM	Retrieving hardware ID from source (172.31.206.200)
9/27/2016 7:55:21 AM	Source (172.31.206.200) hardware ID: 7c536f80-8b73-4d85-aa26-0783aeca7a94
9/27/2016 7:55:21 AM	Retrieving hardware ID from target (localhost)
9/27/2016 7:55:21 AM	Target (localhost) hardware ID: 3bb28b90-ce52-440a-a055-019b612b7712
9/27/2016 7:55:21 AM	Hardware IDs as follows: Source = '7c536f80-8b73-4d85-aa26-0783aeca7a94', Target = '3...
9/27/2016 7:55:21 AM	Completing initialization of new job 7c064253-4cda-495c-a064-29be5aabdf41 (ALPHA to B...
9/27/2016 7:55:21 AM	Initialization of job 7c064253-4cda-495c-a064-29be5aabdf41 (ALPHA to BETA) complete
9/27/2016 7:55:21 AM	Changing to StoppedState from UninitializedState in response to InitializeEvent consumed...
9/27/2016 7:55:21 AM	Exited UninitializedState
9/27/2016 7:55:21 AM	Entered InitializedState
9/27/2016 7:55:21 AM	Starting monitor 7c931961-bca7-48c0-89a0-b972c2679b8c: name = FilesAndFolders_7c06...
9/27/2016 7:55:21 AM	Entered StoppedState
9/27/2016 7:55:21 AM	Stopping monitor 7c931961-bca7-48c0-89a0-b972c2679b8c: Name = FilesAndFolders_7c0...
9/27/2016 7:55:21 AM	Stopping share monitoring
9/27/2016 7:55:22 AM	Changing connection health to Warning
9/27/2016 7:55:22 AM	Event log entry written: '6008'.
9/27/2016 7:55:22 AM	Deleted replication set named FilesAndFolders_7c0642534cda495ca06429be5aabdf41
9/27/2016 7:55:22 AM	Successfully created connection fec22952-ef1c-40df-951a-9cfa7bdf9c8d connecting FilesAn...
9/27/2016 7:55:22 AM	Attaching to engine connection on 172.31.206.200:6325 with following criteria:Guid = fec...
9/27/2016 7:55:22 AM	Waiting 00:10:00 for source endpoint of 'FilesAndFolders_7c0642534cda495ca06429be5aa...
9/27/2016 7:55:23 AM	Established source endpoint of '172.31.206.200:6320' for engine connection with replicatio...
9/27/2016 7:55:23 AM	Updating failover options
9/27/2016 7:55:25 AM	The Double-Take engine is initialized.
9/27/2016 7:55:25 AM	Double-Take is NOT licensed to monitor or assume the identity of another machine.
9/27/2016 7:55:25 AM	The Double-Take engine source module is initialized.
9/27/2016 7:55:25 AM	The Double-Take engine target module is initialized.
9/27/2016 7:55:28 AM	Starting share monitoring
9/27/2016 7:55:32 AM	Changing targetActivationCode health to Ok
9/27/2016 7:55:32 AM	Changing sourceActivationCode health to Ok
9/27/2016 7:55:32 AM	Event log entry written: '6004'.
9/27/2016 7:55:32 AM	Changing to ConnectedState from StoppedState in response to StartSucceededEvent consu...
9/27/2016 7:55:32 AM	Exited StoppedState
9/27/2016 7:55:32 AM	Entered ConnectedState
9/27/2016 7:55:32 AM	Subscribing to engine connection.
9/27/2016 7:55:32 AM	Changing connection health to Ok
9/27/2016 7:55:32 AM	Changing to SynchronizedState from ConnectedState in response to MirrorCompletedEvent...
9/27/2016 7:55:32 AM	Entered ProtectingState
9/27/2016 7:55:32 AM	Entered SynchronizedState
9/27/2016 7:55:32 AM	Persisting shares

The following table identifies the controls and the table columns in the **Job logs** window.

Start

This button starts the addition and scrolling of new messages in the window.

Pause

This button pauses the addition and scrolling of new messages in the window. This is only for the **Job logs** window. The messages are still logged to their respective files on the server.

Copy

This button copies the messages selected in the **Job logs** window to the Windows clipboard.

Clear

This button clears the **Job logs** window. The messages are not cleared from the respective files on the server. If you want to view all of the messages again, close and reopen the **Job logs** window.

Time

This column in the table indicates the date and time when the message was logged.

Description

This column in the table displays the actual message that was logged.

Chapter 6 Recovery


If necessary, create a recovery job to recover the data or a server protected by DR protection job.

- *Creating a recovery job* on page 108—This section includes step-by-step instructions for creating a recovery job.
- *Managing and controlling recovery jobs* on page 130—You can view overview status information about your recovery job.
 - *Viewing recovery job details* on page 140—You can view detailed status information and statistics for your recovery job.
 - *Validating a recovery job* on page 144—Over time, you may want to confirm that any changes in your network or environment have not impacted your job. Use these instructions to validate an existing job.
 - *Viewing a recovery job log* on page 145—You can view the log associated with a specific job.

Creating a recovery job

Use the following instructions to create a recovery job.


1. Before you begin, make sure you have Double-Take installed and licensed on the recovery server.
2. You have two options for creating a recovery job.
 - **From target repository server**—If you do not have an existing protection job, right-click your target repository server on the **Servers** page and select **Recover Data**.
 - **From protection job**—If you have an existing protection job, right-click the protection job on the **Jobs** page and select **Recover Data**.
3. Choose the image that you want to recover. When you select a source in the image list, the snapshots and virtual disks associated with that source will be displayed. Select if you want to recover the live data or the snapshot data. If you select to recover to snapshot data, be sure and select the snapshot.



4. Click **Next** to continue.
5. Choose the data from your source image that you want to recover.
 - **Full Server DR Recovery**—Select this option if you want to recover your entire server. You must have protected your entire server to be able to recover your entire server. If you only protected data, this option will not be available.

- **Files and Folders DR Recovery**—Select this option if you only want to recover data. You can select this option and recover specific data, even if you originally protected your entire server.

If the workload you are looking for is not displayed, enable **Show all workload types**. The workload types in gray text are not available for the source server you have selected. Hover your mouse over an unavailable workload type to see a reason why this workload type is unavailable for the selected source.



6. To select your files and folders in more detail, click the **Replication Rules** heading and expand the volumes under **Folders**.



If you return to this page using the **Back** button in the job creation workflow, your **Workload Types** selection will be rebuilt, potentially overwriting any manual replication rules that you specified. If you do return to this page, confirm your **Workload Types** and **Replication Rules** are set to your desired settings before proceeding forward again.

7. Click **Next** to continue.
8. Choose your target recovery server. This is the server where the protected data will be recovered.



- **Current Servers**—This list contains the servers currently available in your console session. Select your target server from the list. Servers that are not licensed for the workflow you have selected and those not applicable to the workload type you have selected will be filtered out of the list. For example, a DR repository server cannot be used as a recovery server if you have selected to recover an entire protected server. If the server you are looking for is not displayed, enable **Show all servers**. The servers in red are not available for the source server or workload type you have selected. Hover your mouse over an unavailable server to see a reason why this server is unavailable.
- **Find a New Server**—If the server you need is not in the **Current Servers** list, click the **Find a New Server** heading. From here, you can specify a server along with credentials for logging in to the server. If necessary, you can click **Browse** to select a server from a network drill-down list.



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

When specifying credentials for a new server, specify a user that is a member of the local Double-Take Admin and local administrator security groups.

9. Click **Next** to continue.



You may be prompted for a route from the target to the source. This route, and a port if you are using a non-default port, is used so the target can communicate with the source to build job options. This dialog box will be displayed, only if needed.

10. You have many options available for your DR recovery job. Some options are available for both

data and full server jobs, while some options are only for one job type or the other. Configure those options that are applicable to your environment and job type.

Go to each page identified below to see the options available for that section of the **Set Options** page. After you have configured your options, continue with the next step on page 127.

- *General* on page 112
- *Recovery Options* on page 113
- *Recovery Identity* on page 114
- *Network Adapter Options* on page 116
- *Mirror, Verify & Orphaned Files* on page 117
- *Network Route* on page 120
- *Path Mapping* on page 121
- *Compression* on page 123
- *Bandwidth* on page 124
- *Scripts* on page 126


General



A screenshot of a software configuration window titled "General". The window has a light gray header bar with the word "General" and a small upward-pointing arrow icon. Below the header, the text "Job name:" is followed by a text input field containing the text "alpha to beta".

For the **Job name**, specify a unique name for your job.

Recovery Options



Recovery Options

Wait for user to initiate 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.

- **Wait for user to initiate recovery**—The recovery process can wait for you to initiate it, allowing you to control when the recovery finishes. When the recovery mirror is complete, the job will wait in the **Protecting** state for you to manually finish the recovery. Disable this option if you want recovery to finish immediately after the mirror is complete.
- **Shutdown source server**—Specify if you want to shut down the source server, if it is still running, before recovery completes. This option prevents identity conflicts on the network in those cases where the source and target are still both running and communicating.

Recovery Identity

Recovery Identity

Apply source network configuration to the target (Recommended for LAN configurations)

Retain target network configuration (Recommended for WAN configurations)

Update DNS server

DNS Options

Credentials for **domain.com**
User name: **administrator**

Change...

These DNS servers will be updated during failover:

112.42.48.9 Remove

Update these source DNS entries with the corresponding target IP address:

Source Address	Target Address
----------------	----------------

Update TTL (seconds):
300



This section is only applicable to full server recovery jobs.

- **Apply source network configuration to the target**—If you select this option, you can configure the source IP addresses to failover to the target. If your target is on the same subnet as the source (typical of a LAN environment), you should select this option. Do not select this option if you are using a NAT environment that has a different subnet on the other side of the router.



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

- **Retain target network configuration**—If you select this option, the target will retain all of its original IP addresses. If your target is on a different subnet (typical of a WAN or

NAT environment), you should select this option.

- **Update DNS server**—Specify if you want Double-Take to update your DNS server on recovery. If DNS updates are made, the DNS records will be locked during recovery.



DNS updates are not available for Server Core servers or source servers that are in a workgroup.

Make sure port 53 is open for DNS protocol from the recovery server to the DNS servers so the recovery server can discover the source DNS records.

If you are updating DNS but your source and target are in a workgroup, the DNS suffix must be specified for the source NICs and that suffix must correspond to the zone name on the DNS server.

Expand the **DNS Options** section to configure how the updates will be made. The DNS information will be discovered and displayed. If your servers are in a workgroup, you must provide the DNS credentials before the DNS information can be discovered and displayed.

- **Change**—If necessary, click this button and specify a user that has privileges to access and modify DNS records. The account must be a member of the DnsAdmins group for the domain, and must have full control permissions on the source's A (host) and PTR (reverse lookup) records. These permissions are not included by default in the DnsAdmins group.
- **Remove**—If there are any DNS servers in the list that you do not want to update, highlight them and click **Remove**.
- **Update these source DNS entries with the corresponding target IP address**—For each IP address on the source, specify what address you want DNS to use after failover.
- **Update TTL**—Specify the length of time, in seconds, for the time to live value for all modified DNS A records. Ideally, you should specify 300 seconds (5 minutes) or less.



DNS updates will be disabled if the target server cannot communicate with both the source and target DNS servers

Network Adapter Options

Network Adapter Options

Map source network adapters to target network adapters:

Source Network Adapter	Target Network Adapter
Local Area Connection	Local Area Connection
Local Area Connection 2	Local Area Connection 2



This section is only applicable to full server recovery jobs.

For **Map source network adapters to target network adapters**, specify how you want the IP addresses associated with each NIC on the source to be mapped to a NIC on the target. Do not mix public and private networks.

Mirror, Verify & Orphaned Files

Mirror, Verify & Orphaned Files

Mirror Options

Choose a comparison method and whether to mirror the entire file or only the bytes that differ in each file.

Compare file attributes and data. Send the attributes and bytes that differ. (Recommended for databases)

Verification Options

Enable scheduled verification

Verify on this interval: 1 Days

Begin immediately

Begin at this time: 9/26/2016 10:21:53 AM

Report and comparison options

Report only

Report and mirror files

Compare file attributes and data

General Options

Calculate size of protected data upon connection

Delete orphaned files

- **Mirror Options**—Choose a comparison method and whether to mirror the entire file or only the bytes that differ in each file.



If you are using a data recovery job and were using a database application on the original source or were protecting a domain controller, do not use the compare file attributes only options unless you know for certain that you need it. With database applications and because domain controllers store their data in a database, it is critical that all files, not just some of the files, are mirrored. In this case, you should compare both the attributes and the data.

- **Do not compare files. Send the entire file.**—Double-Take will not perform any comparisons between the files on the source and target. All files will be mirrored to the target, sending the entire file. This is equivalent to selecting the mirror all files option prior to Double-Take version 7.1.
- **Compare file attributes. Send the entire file.**—Double-Take will compare file attributes and will mirror those files that have different attributes, sending the entire file. This is similar to selecting the mirror different files and the only if source is newer options prior to Double-Take version 7.1. This option is not available for full server recovery jobs.
- **Compare file attributes. Send the attributes and bytes that differ.**—Double-Take will compare file attributes and will mirror only the attributes and bytes that are different. This is equivalent to selecting the mirror different files, only if source is newer, and use block checksum options prior to Double-Take version 7.1. This option is not available for full server recovery jobs.
- **Compare file attributes and data. Send the attributes and bytes that differ.**—Double-Take will compare file attributes and the file data and will mirror only the attributes and bytes that are different. This is equivalent to selecting the mirror different files and use block checksum options prior to Double-Take version 7.1. If you are using a database application on your source, select this option.

- **Verification Options**—Choose if you want to periodically confirm that the source replica data on the target is identical to the actual data on the source. Verification creates a log file detailing what was verified as well as which files are not synchronized. If the data is not the same, you can automatically initiate a remirror, if configured. The remirror ensures data integrity between the source and target.
-



Because of the way the Windows Cache Manager handles memory, machines that are doing minimal or light processing may have file operations that remain in the cache until additional operations flush them out. This may make Double-Take files on the target appear as if they are not synchronized. When the Windows Cache Manager releases the operations in the cache on the source and target, the files will be updated on the target.

- **Enable scheduled verification**—When this option is enabled, Double-Take will verify the source replica data on the target.
 - **Verify on this interval**—Specify the interval between verification processes.
 - **Begin immediately**—Select this option if you want to start the verification schedule immediately after the job is established.
 - **Begin at this time**—Select this option if you want to start the verification schedule at the specified date and time.
 - **Report only**—Select this option if you only want to generate a verification report. With this option, no data that is found to be different will be mirrored to the target. Choose how you want the verification to compare the files.
 - **Report and mirror files**—Select this option if you want to generate a verification report and mirror data that is different to the target. Select the comparison method and type of mirroring you want to use. See the previous mirroring methods described under *Mirror Options*.
-



If you are using SQL to create snapshots of a SQL database, the verification report will report the file size of the snapshot files on the source and target as different. This is a reporting issue only. The snapshot file is mirrored and replicated completely to the target.

If you are using HP StorageWorks File Migration Agent, migrated files will incorrectly report modified time stamp differences in the verification report. This is a reporting issue only.

- **General Options**—Choose your general mirroring options. The calculate size option is only available for data protection and recovery. It is not available for full server protection and recovery.
 - **Calculate size of protected data upon connection**—Specify if you want Double-Take to determine the mirroring percentage calculation based on the amount of data being protected. If you enable this option, the calculation will begin when mirroring begins. For the initial mirror, the percentage will display after the calculation is complete, adjusting to the amount of the mirror that has completed during the time it

took to complete the calculation. Subsequent mirrors will initially use the last calculated size and display an approximate percentage. Once the calculation is complete, the percentage will automatically adjust down or up to indicate the amount that has been completed. Disabling calculation will result in the mirror status not showing the percentage complete or the number of bytes remaining to be mirrored.



The calculated amount of protected data may be slightly off if your data set contains compressed or sparse files.

This option is not available for full server recovery jobs.

- **Delete orphaned files**—An orphaned file is a file that exists in the replica data on the target, but does not exist in the protected data on the source. This option specifies if orphaned files should be deleted on the target.
-



Orphaned file configuration is a per target configuration. All jobs to the same target will have the same orphaned file configuration.

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

The orphaned file feature does not delete alternate data streams. To do this, use a full mirror, which will delete the additional streams when the file is re-created.

If you want to move orphaned files rather than delete them, you can configure this option along with the move deleted files feature to move your orphaned files to the specified deleted files directory. See *Target server properties* on page 54 for more information.

During a mirror, orphaned file processing success messages will be logged to a separate orphaned file log on the source. This keeps the Double-Take log from being overrun with orphaned file success processing messages. Orphaned files processing statistics and any errors in orphaned file processing will still be logged to the Double-Take log, and during difference mirrors, verifications, and restorations, all orphaned file processing messages are logged to the Double-Take log. The orphaned file log is located in the **Logging folder** specified for the source. See *Log file properties* on page 59 for details on the location of that folder. The orphaned log file is appended to during each orphaned file processing during a mirror, and the log file will be a maximum of 50 MB.

Network Route



Network Route

Send data to the target server using this route:

10.10.10.30




This section is only applicable for data recovery jobs.

By default, Double-Take will select a target route for transmissions. If desired, specify an alternate route on the target that the data will be transmitted through. This allows you to select a different route for Double-Take traffic. For example, you can separate regular network traffic and Double-Take traffic on a machine with multiple IP addresses. You can also select or manually enter a public IP address (which is the public IP address of the server's router) if you are using a NAT environment. If you enter a public IP addresses, you will see additional fields allowing you to disable the default communication ports and specify other port numbers to use, allowing the target to communicate through a router. The **Management Service port** is used to persist the source share configuration when shares are being protected. The **Replication Service port** is used for data transmission.



If you change the IP address on the target which is used for the target route, you will be unable to edit the job. If you need to make any modifications to the job, it will have to be deleted and re-created.

Path Mapping



This section is only applicable to data recovery jobs.

- **Mappings**—Specify the location on the recovery server where the source data will be restored. By default, the data will be restored in the same directory structure as the original source, in a one to one configuration.
 - **All To One**—Click this button to set the mapping so that the replica source data will be stored on a single volume on the target. The pre-defined path is `\source_name\volume_name`. If you are protecting multiple volumes on the source, each volume would be stored on the same volume on the target. For example, `C:\data` and `D:\files` for the source Alpha would be stored in `D:\alpha\C` and `D:\alpha\D`, respectively.
 - **One To One**—Click this button to set the mapping so that the replica source data will be stored in the same directory structure on the target. For example, `C:\data` and `D:\files` on the source will be stored in `C:\data` and `D:\files`, respectively, on the target.
 - **Custom Location**—If the pre-defined options do not store the data in a location that is appropriate for your network operations, you can specify your own custom location where the replica source data will be stored. Click the **Target Path** and edit it, selecting the appropriate location.



If you are protecting system state data (like your Program Files or Documents and Settings directory), you must select the **All to One** mapping or specify a customized location in order to avoid sharing violations.

If you are protecting dynamic volumes or mount points, your location on the target must be able to accommodate the amount of data that may be stored on the source.


If you are protecting multiple mount points, your directory mapping must not create a cycle or loop. For example, if you have the C: volume mounted at D:\C and the D: volume mounted at C:\D, this is a circular configuration. If you establish a job for either C:\D or D:\C, there will be a circular configuration and Double-Take mirroring will never complete.

If you are protecting sparse files, the amount of disk space available must be equal to or greater than the entire size of the sparse file. If the target location is an NTFS 5

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

- **Block target paths upon connection**—You can block writing to the recovery data on the target. This keeps the data from being changed outside of Double-Take processing.

Compression



To help reduce the amount of bandwidth needed to transmit Double-Take data, compression allows you to compress data prior to transmitting it across the network. In a WAN environment this provides optimal use of your network resources. If compression is enabled, the data is compressed before it is transmitted from the source. When the target receives the compressed data, it decompresses it and then writes it to disk. You can set the level from **Minimum** to **Maximum** to suit your needs.


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

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



All jobs from a single source connected to the same IP address on a target will share the same compression configuration.

Bandwidth



Bandwidth

Do not limit bandwidth

Use a fixed limit (bytes per second):

Preset bandwidth: Bandwidth (bytes per second):

Use scheduled limits:

Days	Start	End	Limit
Every day	6:00 AM	6:00 PM	1250000
Every day	6:00 PM	6:00 AM	5625000

 Bandwidth is not limited during unscheduled times.

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



All jobs from a single source connected to the same IP address on a target will share the same bandwidth configuration.

- **Do not limit bandwidth**—Double-Take will transmit data using 100% bandwidth availability.
- **Use a fixed limit**—Double-Take will transmit data using a limited, fixed bandwidth. Select a **Preset bandwidth** limit rate from the common bandwidth limit values. The **Bandwidth** field will automatically update to the bytes per second value for your selected bandwidth. This is the maximum amount of data that will be transmitted per second. If desired, modify the bandwidth using a bytes per second value. The minimum limit should be 3500 bytes per second.
- **Use scheduled limits**—Double-Take will transmit data using a dynamic bandwidth based on the schedule you configure. Bandwidth will not be limited during unscheduled times.
 - **New**—Click **New** to create a new scheduled bandwidth limit. Specify the following information.
 - **Daytime entry**—Select this option if the start and end times of the bandwidth window occur in the same day (between 12:01 AM and midnight). The start time must occur before the end time.
 - **Overnight entry**—Select this option if the bandwidth window begins on one day and continues past midnight into the next day. The start time must be later than the end time, for example 6 PM to 6 AM.
 - **Day**—Enter the day on which the bandwidth limiting should occur. You can pick a specific day of the week, **Weekdays** to have the limiting occur Monday through Friday, **Weekends** to have the limiting occur Saturday and Sunday, or **Every day** to have the limiting repeat on all days of the week.
 - **Start time**—Enter the time to begin bandwidth limiting.

- **End time**—Enter the time to end bandwidth limiting.
 - **Preset bandwidth**—Select a bandwidth limit rate from the common bandwidth limit values. The **Bandwidth** field will automatically update to the bytes per second value for your select bandwidth.
 - **Bandwidth**—If desired, modify the bandwidth using a bytes per second value. The minimum limit should be 3500 bytes per second.
 - **Edit**—Click **Edit** to modify an existing scheduled bandwidth limit.
 - **Delete**—Click **Delete** to remove a scheduled bandwidth limit.
-



If you change your job option from **Use scheduled limits** to **Do not limit bandwidth** or **Use a fixed limit**, any schedule that you created will be preserved. That schedule will be reused if you change your job option back to **Use scheduled limits**.

You can manually override a schedule after a job is established by selecting **Other Job Options, Set Bandwidth**. If you select **No bandwidth limit** or **Fixed bandwidth limit**, that manual override will be used until you go back to your schedule by selecting **Other Job Options, Set Bandwidth, Scheduled bandwidth limit**. For example, if your job is configured to use a daytime limit, you would be limited during the day, but not at night. But if you override that, your override setting will continue both day and night, until you go back to your schedule. See the *Managing and controlling jobs* section for your job type for more information on the **Other Job Options**.

Scripts

The screenshot shows a window titled "Scripts" with three sections:

- Mirror Start**: Script file: `c:\scripts\mirrorstart.bat`, Arguments: (empty), Allow script to interact with desktop, Delay until script completes.
- Mirror Complete**: Script file: (empty), Arguments: (empty), Allow script to interact with desktop, Delay until script completes.
- Mirror Stop**: Script file: `c:\scripts\mirrorcomplete.bat`, Arguments: `arg1`, Allow script to interact with desktop, Delay until script completes.

You can customize mirroring by running scripts on the target at pre-defined points in the mirroring process. Scripts may contain any valid Windows command, executable, or batch file. The scripts are processed using the same account running the Double-Take service, unless you have identified a specific account through the server's properties. See *Script credentials* on page 58. There are three types of mirroring scripts.

- **Mirror Start**—This script starts when the target receives the first mirror operation. In the case of a difference mirror, this may be a long time after the mirror is started because the script does not start until the first different data is received on the target. If the data is synchronized and a difference mirror finds nothing to mirror, the script will not be executed. Specify the full path and name of the **Script file**.
- **Mirror Complete**—This script starts when a mirror is completed. Because the mirror statistics may indicate a mirror is at 99-100% when it is actually still processing (for example, if files were added after the job size was calculated, if there are alternate data streams, and so on), the script will not start until all of the mirror data has been completely processed on the target. Specify the full path and name of the **Script file**.
- **Mirror Stop**—This script starts when a mirror is stopped, which may be caused by an auto-disconnect occurring while a mirror is running, the service is shutdown while a mirror is running, or if you stop a mirror manually. Specify the full path and name of the **Script file**.
- **Arguments**—Specify a comma-separated list of valid arguments required to execute the script.
- **Allow script to interact with desktop**—Enable this option if you want the script processing to be displayed on the screen. Otherwise, the script will execute silently in the background. If you enable this option, the account running the script must have permission to interact with the desktop.
- **Delay until script completes**—Enable this option if you want to delay the mirroring process until the associated script has completed. If you select this option, make sure your script handles errors, otherwise the mirroring process may never complete if the process is waiting on a script that cannot complete.
- **Test**—You can test your script manually by clicking **Test**. Your script will be executed if you test it. If necessary, manually undo any changes that you do not want on your target after testing the script.



If you establish mirroring scripts for one job and then establish additional jobs to the same target using the same target path mapping, the mirroring scripts will automatically be applied to those subsequent jobs. If you select a different target path mapping, the mirroring scripts will have to be reconfigured for the new job(s).

11. Click **Next** to continue.
12. Double-Take validates that your source and target are compatible. The **Summary** page displays your options and validation items.

Errors are designated by a white X inside a red circle. Warnings are designated by a black exclamation point (!) inside a yellow triangle. A successful validation is designated by a white checkmark inside a green circle. You can sort the list by the icon to see errors, warnings, or successful validations together. Click on any of the validation items to see details. You must correct any errors before you can continue. Depending on the error, you may be able to click **Fix** or **Fix All** and let Double-Take correct the problem for you. For those errors that Double-Take cannot correct automatically, you will need to modify the source or target to correct the error, or you can select a different target. You must revalidate the selected servers, by clicking **Recheck**, until the validation check passes without errors.

If you receive a path transformation error during job validation indicating a volume does not exist on the target server, even though there is no corresponding data being protected on the source, you will need to manually modify your replication rules. Go back to the **Choose Data** page and under the **Replication Rules**, locate the volume from the error message. Remove any rules associated with that volume. Complete the rest of the workflow and the validation should pass.

Before a job is created, the results of the validation checks are logged to the Double-Take Management Service log on the target. After a job is created, the results of the validation checks are logged to the job log. See the Double-Take *Reference Guide* for details on the various Double-Take log files.

13. Once your servers have passed validation and you are ready to establish recovery, click **Finish**, and you will automatically be taken to the **Jobs** page.
-



Jobs in a NAT environment may take longer to start.

14. When the recovery mirror has completed, the job may or may not terminate automatically depending on your selection for **Wait for user intervention before recovery**. If you disabled user intervention, the job will automatically terminate to complete the recovery process. If you enabled user intervention, when the recovery mirror is complete, the status will change to **Protecting**. Use this time to complete any necessary tasks. When you are ready to complete the recovery, highlight the job on the **Jobs** page and click **Failover, Cutover, or Recover**.
15. Select how you want to handle any data in the target queue.
 - **Apply data in target queues before recovery**—All of the data in the target queue will be applied before recovery finishes. The advantage to this option is that all of the data that the target has received will be applied before recovery completes. The disadvantage to this option is depending on the amount of data in queue, the amount of time to apply all of the data could be lengthy.

- **Discard data in the target queues and recover immediately**—All of the data in the target queue will be discarded and recovery will be completed immediately. The advantage to this option is that recover will complete immediately. The disadvantage is that any data in the target queue will be lost.

16. Click **Recover** to complete the recovery process.



If you are recovering an entire server, keep in the mind the following notes.

- If your NICs were configured for network load balancing (NLB), you will have to reconfigure that after recovery.
- Right before recovery occurs, Double-Take will stop all services that are not critical to Windows. If the stop command fails (perhaps because it is a blocking driver that cannot be shutdown, as is the case with some anti-virus software) or a third-party tool restarts any of these services, Double-Take may not be able to successfully recover files locked by the services. In this case, you may have to make manual modifications to your server after recovery.
- Some applications and hardware devices create and use software devices within the operating system, but they have the characteristics of a hardware device. For example, NIC teaming solutions are typically implemented in the operating system, however they are still designed to emulate a single piece of network hardware. In these cases, the device will not be recovered because it appears to be a hardware device.
- Because the Windows product activation is dependent on hardware, you may need to reactivate your Windows registration after recovery. The reactivation depends on several factors including service pack level, Windows edition, and your licensing type. If a target comes online after recovery with an activation failure, use the steps below appropriate for your license type. Additionally, if you are using Windows 2012, you may only have 60 minutes to complete the reactivation process until Windows activation tampering automatically shuts down your server.
 - **Retail licensing**—Retail licensing allows the activation of a single operating system installation.
 1. Open the **System** applet in Windows **Control Panel**.
 2. Under **Windows activation** at the bottom of the page, click **Change product key**.
 3. Enter your retail license key. You may need access to the Internet or to call Microsoft to complete the activation.
 - **MAK volume licensing**—Multiple Activation Key (MAK) licensing allows the activation of multiple operating system installations using the same activation key.
 1. View or download the Microsoft Volume Activation Deployment Guide from the Microsoft web site.
 2. Using an administrative user account, open a command prompt and follow the instructions from the deployment guide to activate MAK clients. Multiple reboots may be necessary before you can access a command prompt. You may need access to the Internet or to call Microsoft to complete the activation.
 - **KMS volume licensing**—Key Management Service (KMS) licensing allows IT professionals to complete activations on their local network without contacting Microsoft.

1. View or download the Microsoft Volume Activation Deployment Guide from the Microsoft web site.
 2. Using an administrative user account, open a command prompt and follow the instructions from the deployment guide to convert a MAK activation client to a KMS client. Multiple reboots may be necessary before you can access a command prompt.
- After a recovery is complete and your recovery server is online as your source, when you login, you may have to specify a Windows reboot reason. You can specify any reason and continue. Additionally, you may see a prompt indicating a reboot is required because of a device change. You can disregard this error and select to reboot later.

Because Windows 64-bit has a strict driver signing policy, if you get a stop code 0x7b after failover, you may have drivers failing to load because the driver signatures are failing the policy. In this case, reboot the server and press F8. Choose the option to not enforce the driver signing policy. If this allows the system to boot, then the problem is being caused by the cat file signature mismatch. If your system still fails to boot, contact technical support.

Managing and controlling recovery jobs

Click **Jobs** from the main Double-Take Console toolbar. The **Jobs** page allows you to view status information about your jobs. You can also control your jobs from this page.

The jobs displayed in the right pane depend on the server group folder selected in the left pane. Every job for each server in your console session is displayed when the **Jobs on All Servers** group is selected. If you have created and populated server groups (see *Managing servers* on page 24), then only the jobs associated with the server or target servers in that server group will be displayed in the right pane.

- *Overview job information displayed in the top right pane* on page 130
- *Detailed job information displayed in the bottom right pane* on page 133
- *Job controls* on page 135

Overview job information displayed in the top right pane

The top pane displays high-level overview information about your jobs. You can sort the data within a column in ascending and descending order. You can also move the columns to the left or right of each other to create your desired column order. The list below shows the columns in their default left to right order.

If you are using server groups, you can filter the jobs displayed in the top right pane by expanding the **Server Groups** heading and selecting a server group.

Column 1 (Blank)

The first blank column indicates the state of the job.

 A green circle with a white checkmark indicates the job is in a healthy state. No action is required.

 A yellow triangle with a black exclamation point indicates the job is in a pending or warning state. This icon is also displayed on any server groups that you have created that contain a job in a pending or warning state. Double-Take is working or waiting on a pending process or attempting to resolve the warning state.

 A red circle with a white X indicates the job is in an error state. This icon is also displayed on any server groups that you have created that contain a job in an error state. You will need to investigate and resolve the error.

 The job is in an unknown state.

Job

The name of the job

Source Server

The name of the source. This could be the name or IP address of your source.

Target Server

The name of the target. This could be the name or IP address of your target.

Job Type

Each job type has a unique job type name. This job is either a Full Server to Windows DR Recovery job or a Data Only for Windows DR Recovery job, depending on if you are recovering the entire source server or just specific data from the source server. For a complete list of all job type names, press F1 to view the Double-Take Console online help.

Activity

There are many different **Activity** messages that keep you informed of the job activity. Most of the activity messages are informational and do not require any administrator interaction. If you see error messages, check the job details. Keep in mind that **Idle** indicates console to server activity is idle, not that your servers are idle.

Mirror Status

- **Calculating**—The amount of data to be mirrored is being calculated.
- **In Progress**—Data is currently being mirrored.
- **Waiting**—Mirroring is complete, but data is still being written to the target.
- **Idle**—Data is not being mirrored.
- **Paused**—Mirroring has been paused.
- **Stopped**—Mirroring has been stopped.
- **Removing Orphans**—Orphan files on the target are being removed or deleted depending on the configuration.
- **Verifying**—Data is being verified between the source and target.
- **Unknown**—The console cannot determine the status.

Replication Status

- **Replicating**—Data is being replicated to the target.
- **Ready**—There is no data to replicate.
- **Pending**—Replication is pending.
- **Stopped**—Replication has been stopped.
- **Out of Memory**—Replication memory has been exhausted.
- **Failed**—The Double-Take service is not receiving replication operations from the Double-Take driver. Check the Event Viewer for driver related issues.
- **Unknown**—The console cannot determine the status.

Transmit Mode

- **Active**—Data is being transmitted to the target.
- **Paused**—Data transmission has been paused.
- **Scheduled**—Data transmission is waiting on schedule criteria.
- **Stopped**—Data is not being transmitted to the target.

- **Error**—There is a transmission error.
- **Unknown**—The console cannot determine the status.

Operating System

The job type operating system

Detailed job information displayed in the bottom right pane

The details displayed in the bottom pane provide additional information for the job highlighted in the top pane. You can expand or collapse the bottom pane by clicking on the **Job Highlights** heading.

Name

The name of the job

Target data state

- **OK**—The data on the target is in a good state.
- **Mirroring**—The target is in the middle of a mirror process. The data will not be in a good state until the mirror is complete.
- **Mirror Required**—The data on the target is not in a good state because a remirror is required. This may be caused by an incomplete or stopped mirror or an operation may have been dropped on the target.
- **Busy**—The source is low on memory causing a delay in getting the state of the data on the target.
- **Not Loaded**—Double-Take target functionality is not loaded on the target server. This may be caused by a license key error.
- **Unknown**—The console cannot determine the status.

Mirror remaining

The total number of mirror bytes that are remaining to be sent from the source to the target. This value may be zero if you have enabled **Mirror only changed files when source reboots** on the *Server setup properties* on page 44.

Mirror skipped

The total number of bytes that have been skipped when performing a difference. These bytes are skipped because the data is not different on the source and target. This value may be zero if you have enabled **Mirror only changed files when source reboots** on the *Server setup properties* on page 44.

Replication queue

The total number of replication bytes in the source queue

Disk queue

The amount of disk space being used to queue data on the source

Recovery point latency

The length of time replication is behind on the target compared to the source. This is the time period of data that would be lost if a failure were to occur at the current time.

Bytes sent

The total number of mirror and replication bytes that have been transmitted to the target

Bytes sent (compressed)

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

Connected since

The date and time indicating when the current job was started.

Recent activity

Displays the most recent activity for the selected job, along with an icon indicating the success or failure of the last initiated activity. Click the link to see a list of recent activities for the selected job. You can highlight an activity in the list to display additional details about the activity.

Image name

The name of the image on the repository server

Additional information

Depending on the current state of your job, you may see additional information displayed to keep you informed about the progress and status of your job. If there is no additional information, you will see (None) displayed.

Job controls

You can control your job through the toolbar buttons available on the **Jobs** page. If you select multiple jobs, some of the controls will apply only to the first selected job, while others will apply to all of the selected jobs. For example, **View Job Details** will only show details for the first selected job, while **Stop** will stop protection for all of the selected jobs.

If you want to control just one job, you can also right click on that job and access the controls from the pop-up menu.

View Job Details

This button leaves the **Jobs** page and opens the **View Job Details** page.

Edit Job Properties

This button leaves the **Jobs** page and opens the **EditJob Properties** page.

Delete

Stops (if running) and deletes the selected jobs.

Provide Credentials

Changes the login credentials that the job (which is on the target machine) uses to authenticate to the servers in the job. This button opens the Provide Credentials dialog box where you can specify the new account information and which servers you want to update. See *Providing server credentials* on page 37. You will remain on the **Jobs** page after updating the server credentials. If your servers use the same credentials, make sure you also update the credentials on the **Servers** page so that the Double-Take Console can authenticate to the servers in the console session. See *Managing servers* on page 24.

View Recent Activity

Displays the recent activity list for the selected job. Highlight an activity in the list to display additional details about the activity.

Start

Starts or resumes the selected jobs.

If you have previously stopped protection, the job will restart mirroring and replication.

If you have previously paused protection, the job will continue mirroring and replication from where it left off, as long as the Double-Take queue was not exhausted during the

time the job was paused. If the Double-Take queue was exhausted during the time the job was paused, the job will restart mirroring and replication.

Also if you have previously paused protection, all jobs from the same source to the same IP address on the target will be resumed.

Pause

Pauses the selected jobs. Data will be queued on the source while the job is paused.

All jobs from the same source to the same IP address on the target will be paused.

Stop

Stops the selected jobs. The jobs remain available in the console, but there will be no mirroring or replication data transmitted from the source to the target. Mirroring and replication data will not be queued on the source while the job is stopped, requiring a remirror when the job is restarted. The type of remirror will depend on your job settings.

Take Snapshot

Snapshots are not applicable to DR recovery jobs.

Manage Snapshots

Snapshots are not applicable to DR recovery jobs.

Failover, Cutover, or Recover

Finishes the recovery process. See the final steps of the process for *Creating a recovery job* on page 108.

Failback

Starts the failback process. Failback does not apply to DR jobs.

Restore

Starts the restoration process. Restore does not apply to DR jobs.

Reverse

Reverses protection. Reverse protection does not apply to DR jobs.

Recover Data

Recovers the selected DR job. You will be taken directly to the **Choose Source Image to Recover** page, which bypasses the first three pages of the recovery workflow (**Get Started**, **Double-Take DR**, and **Choose Source Repository**

Server), automatically using the repository server of the selected protection job. See *Creating a recovery job* on page 108 for details on recovering a DR job.

Undo Failover or Cutover

Cancels a failover or cutover by undoing it. Undo failover does not apply to DR jobs.

View Job Log

Opens the job log. On the right-click menu, this option is called **View Logs**, and you have the option of opening the job log, source server log, or target server log.

Other Job Actions

Opens a small menu of other job actions. These job actions will be started immediately, but keep in mind that if you stop and restart your job, the job's configured settings will override any other job actions you may have initiated.

- **Mirroring**—You can start, stop, pause and resume mirroring for any job that is running.

When pausing a mirror, Double-Take stops queuing mirror data on the source but maintains a pointer to determine what information still needs to be mirrored to the target. Therefore, when resuming a paused mirror, the process continues where it left off.

When stopping a mirror, Double-Take stops queuing mirror data on the source and does not maintain a pointer to determine what information still needs to be mirrored to the target. Therefore, when starting a mirror that has been stopped, you will need to decide what type of mirror to perform.

- **Mirror Options**—Choose a comparison method and whether to mirror the entire file or only the bytes that differ in each file.



If you are using a data recovery job and were using a database application on the original source or were protecting a domain controller, do not use the compare file attributes only options unless you know for certain that you need it. With database applications and because domain controllers store their data in a database, it is critical that all files, not just some of the files, are mirrored. In this case, you should compare both the attributes and the data.

-
- **Do not compare files. Send the entire file.**—Double-Take will not perform any comparisons between the files on the source and target. All files will be mirrored to the target, sending the entire file. This is equivalent to selecting the mirror all files option prior to Double-Take version 7.1.
 - **Compare file attributes. Send the entire file.**—Double-Take will compare file attributes and will mirror those files that have different

attributes, sending the entire file. This is similar to selecting the mirror different files and the only if source is newer options prior to Double-Take version 7.1. This option is not available for full server recovery jobs.

- **Compare file attributes. Send the attributes and bytes that differ.**—Double-Take will compare file attributes and will mirror only the attributes and bytes that are different. This is equivalent to selecting the mirror different files, only if source is newer, and use block checksum options prior to Double-Take version 7.1. This option is not available for full server recovery jobs.
- **Compare file attributes and data. Send the attributes and bytes that differ.**—Double-Take will compare file attributes and the file data and will mirror only the attributes and bytes that are different. This is equivalent to selecting the mirror different files and use block checksum options prior to Double-Take version 7.1. If you are using a database application on your source, select this option.
- **Calculate size of protected data before mirroring**—Specify if you want Double-Take to determine the mirroring percentage calculation based on the amount of data being protected. If you enable this option, the calculation will begin when mirroring begins. For the initial mirror, the percentage will display after the calculation is complete, adjusting to the amount of the mirror that has completed during the time it took to complete the calculation. Subsequent mirrors will initially use the last calculated size and display an approximate percentage. Once the calculation is complete, the percentage will automatically adjust down or up to indicate the amount that has been completed. Disabling calculation will result in the mirror status not showing the percentage complete or the number of bytes remaining to be mirrored.



The calculated amount of protected data may be slightly off if your data set contains compressed or sparse files.

Do not disable this option for full server jobs. The calculation time is when the source generates a hard link report that the target will use to process hard links on the target. If you disable the calculation and thus the report generation, the hard link processing that occurs on the target will be ineffective and you may have problems after failover.

-
- **Verify**—Even if you have scheduled the verification process, you can run it manually any time a mirror is not in progress.
 - **Report only**—Select this option if you only want to generate a verification report. With this option, no data that is found to be different will be mirrored to the target. Choose how you want the verification to compare the files.
 - **Report and mirror files**—Select this option if you want to generate a verification report and mirror data that is different to the target. Select the comparison method and type of mirroring you want to use. See the previous mirroring methods described under *Mirror Options*.
 - **Set Bandwidth**—You can manually override bandwidth limiting settings configured for your job at any time.

- **No bandwidth limit**—Double-Take will transmit data using 100% bandwidth availability.
- **Fixed bandwidth limit**—Double-Take will transmit data using a limited, fixed bandwidth. Select a **Preset bandwidth** limit rate from the common bandwidth limit values. The **Bandwidth** field will automatically update to the bytes per second value for your selected bandwidth. This is the maximum amount of data that will be transmitted per second. If desired, modify the bandwidth using a bytes per second value. The minimum limit should be 3500 bytes per second.
- **Scheduled bandwidth limit**—If your job has a configured scheduled bandwidth limit, you can enable that schedule with this option.
- **Delete Orphans**—Even if you have enabled orphan file removal during your mirror and verification processes, you can manually remove them at any time.
- **Target**—You can pause the target, which queues any incoming Double-Take data from the source on the target. All active jobs to that target will complete the operations already in progress. Any new operations will be queued on the target until the target is resumed. The data will not be committed until the target is resumed. Pausing the target only pauses Double-Take processing, not the entire server.

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

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

Filter

Select a filter option from the drop-down list to only display certain jobs. You can display **Healthy jobs**, **Jobs with warnings**, or **Jobs with errors**. To clear the filter, select **All jobs**. If you have created and populated server groups, then the filter will only apply to the jobs associated with the server or target servers in that server group. See *Managing servers* on page 24.

Search

Allows you to search the source or target server name for items in the list that match the criteria you have entered.

Overflow Chevron

Displays any toolbar buttons that are hidden from view when the window size is reduced.

Viewing recovery job details

From the **Jobs** page, highlight the job and click **View Job Details** in the toolbar.

Review the following table to understand the detailed information about your job displayed on the **View Job Details** page.

Job name

The name of the job

Job type

Each job type has a unique job type name. This job is either a Full Server to Windows DR Recovery job or a Data Only for Windows DR Recovery job, depending on if you are recovering the entire source server or just specific data from the source server. For a complete list of all job type names, press F1 to view the Double-Take Console online help.

Health

-  The job is in a healthy state.
-  The job is in a warning state.
-  The job is in an error state.
-  The job is in an unknown state.

Activity

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

Connection ID

The incremental counter used to number connections. The number is incremented when a connection is created. The counter is reset if there are no existing jobs and the Double-Take service is restarted.

Transmit mode

- **Active**—Data is being transmitted to the target.
- **Paused**—Data transmission has been paused.
- **Scheduled**—Data transmission is waiting on schedule criteria.
- **Stopped**—Data is not being transmitted to the target.
- **Error**—There is a transmission error.
- **Unknown**—The console cannot determine the status.

Target data state

- **OK**—The data on the target is in a good state.
- **Mirroring**—The target is in the middle of a mirror process. The data will not be in a good state until the mirror is complete.
- **Mirror Required**—The data on the target is not in a good state because a remirror is required. This may be caused by an incomplete or stopped mirror or an operation may have been dropped on the target.
- **Busy**—The source is low on memory causing a delay in getting the state of the data on the target.
- **Not Loaded**—Double-Take target functionality is not loaded on the target server. This may be caused by a license key error.
- **Unknown**—The console cannot determine the status.

Target route

The IP address on the target used for Double-Take transmissions.

Compression

- **On / Level**—Data is compressed at the level specified.
- **Off**—Data is not compressed.

Encryption

- **On**—Data is being encrypted before it is sent from the source to the target.
- **Off**—Data is not being encrypted before it is sent from the source to the target.

Bandwidth limit

If bandwidth limiting has been set, this statistic identifies the limit. The keyword **Unlimited** means there is no bandwidth limit set for the job.

Connected since

The source server date and time indicating when the current job was started. This field is blank, indicating that a TCP/IP socket is not present, when the job is waiting on transmit options or if the transmission has been stopped. This field will maintain the date and time, indicating that a TCP/IP socket is present, when transmission has been paused.

Additional information

Depending on the current state of your job, you may see additional information displayed to keep you informed about the progress and status of your job. If there is no additional information, you will see (None) displayed.

Mirror status

- **Calculating**—The amount of data to be mirrored is being calculated.
- **In Progress**—Data is currently being mirrored.
- **Waiting**—Mirroring is complete, but data is still being written to the target.
- **Idle**—Data is not being mirrored.

- **Paused**—Mirroring has been paused.
- **Stopped**—Mirroring has been stopped.
- **Removing Orphans**—Orphan files on the target are being removed or deleted depending on the configuration.
- **Verifying**—Data is being verified between the source and target.
- **Unknown**—The console cannot determine the status.

Mirror percent complete

The percentage of the mirror that has been completed

Mirror remaining

The total number of mirror bytes that are remaining to be sent from the source to the target. This value may be zero if you have enabled **Mirror only changed files when source reboots** on the *Server setup properties* on page 44.

Mirror skipped

The total number of bytes that have been skipped when performing a difference. These bytes are skipped because the data is not different on the source and target. This value may be zero if you have enabled **Mirror only changed files when source reboots** on the *Server setup properties* on page 44.

Replication status

- **Replicating**—Data is being replicated to the target.
- **Ready**—There is no data to replicate.
- **Pending**—Replication is pending.
- **Stopped**—Replication has been stopped.
- **Out of Memory**—Replication memory has been exhausted.
- **Failed**—The Double-Take service is not receiving replication operations from the Double-Take driver. Check the Event Viewer for driver related issues.
- **Unknown**—The console cannot determine the status.

Replication queue

The total number of replication bytes in the source queue

Disk queue

The amount of disk space being used to queue data on the source

Bytes sent

The total number of mirror and replication bytes that have been transmitted to the target

Bytes sent compressed

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

Recovery point latency

The length of time replication is behind on the target compared to the source. This is the time period of data that would be lost if a failure were to occur at the current time.

Mirror start time

The UTC time when mirroring started

Mirror end time

The UTC time when mirroring ended

Total time for last mirror

The length of time it took to complete the last mirror process

Validating a recovery job

Over time, you may want to confirm that any changes in your network or environment have not impacted your Double-Take job. Use these instructions to validate an existing job.


1. From the **Jobs** page, highlight the job and click **View Job Details** in the toolbar.
2. In the **Tasks** area on the right on the **View Job Details** page, click **Validate job properties**.
3. Errors are designated by a white X inside a red circle. Warnings are designated by a black exclamation point (!) inside a yellow triangle. A successful validation is designated by a white checkmark inside a green circle. You can sort the list by the icon to see errors, warnings, or successful validations together. Click on any of the validation items to see details. You must correct any errors before you can continue. Depending on the error, you may be able to click **Fix** or **Fix All** and let Double-Take correct the problem for you. For those errors that Double-Take cannot correct automatically, you will need to modify the source or target to correct the error, or you can select a different target. You must revalidate the selected servers, by clicking **Recheck**, until the validation check passes without errors.

Validation checks for an existing job are logged to the job log on the target server.

4. Once your servers have passed validation, click **Close**.

Viewing a recovery job log

You can view a job log file through the Double-Take Console by selecting **View Job Log** from the toolbar on the **Jobs** page. Separate logging windows allow you to continue working in the Double-Take Console while monitoring log messages. You can open multiple logging windows for multiple jobs. When the Double-Take Console is closed, all logging windows will automatically close.



Time	Description
9/27/2016 7:55:21 AM	Retrieving hardware ID from source (172.31.206.200)
9/27/2016 7:55:21 AM	Source (172.31.206.200) hardware ID: 7c536f80-8b73-4d85-aa26-0783aeca7a94
9/27/2016 7:55:21 AM	Retrieving hardware ID from target (localhost)
9/27/2016 7:55:21 AM	Target (localhost) hardware ID: 3bb28b90-ce52-440a-a055-019b612b7712
9/27/2016 7:55:21 AM	Hardware IDs as follows: Source = '7c536f80-8b73-4d85-aa26-0783aeca7a94', Target = '3...
9/27/2016 7:55:21 AM	Completing initialization of new job 7c064253-4cda-495c-a064-29be5aabdf41 (ALPHA to B...
9/27/2016 7:55:21 AM	Initialization of job 7c064253-4cda-495c-a064-29be5aabdf41 (ALPHA to BETA) complete
9/27/2016 7:55:21 AM	Changing to StoppedState from UninitializedState in response to InitializeEvent consumed...
9/27/2016 7:55:21 AM	Exited UninitializedState
9/27/2016 7:55:21 AM	Entered InitializedState
9/27/2016 7:55:21 AM	Starting monitor 7c931961-bca7-48c0-89a0-b972c2679b8c: name = FilesAndFolders_7c06...
9/27/2016 7:55:21 AM	Entered StoppedState
9/27/2016 7:55:21 AM	Stopping monitor 7c931961-bca7-48c0-89a0-b972c2679b8c: Name = FilesAndFolders_7c0...
9/27/2016 7:55:21 AM	Stopping share monitoring
9/27/2016 7:55:22 AM	Changing connection health to Warning
9/27/2016 7:55:22 AM	Event log entry written: '6008'.
9/27/2016 7:55:22 AM	Deleted replication set named FilesAndFolders_7c0642534cda495ca06429be5aabdf41
9/27/2016 7:55:22 AM	Successfully created connection fec22952-ef1c-40df-951a-9cfa7bdf9c8d connecting FilesAn...
9/27/2016 7:55:22 AM	Attaching to engine connection on 172.31.206.200:6325 with following criteria:Guid = fec...
9/27/2016 7:55:22 AM	Waiting 00:10:00 for source endpoint of 'FilesAndFolders_7c0642534cda495ca06429be5aa...
9/27/2016 7:55:23 AM	Established source endpoint of '172.31.206.200:6320' for engine connection with replicatio...
9/27/2016 7:55:23 AM	Updating failover options
9/27/2016 7:55:25 AM	The Double-Take engine is initialized.
9/27/2016 7:55:25 AM	Double-Take is NOT licensed to monitor or assume the identity of another machine.
9/27/2016 7:55:25 AM	The Double-Take engine source module is initialized.
9/27/2016 7:55:25 AM	The Double-Take engine target module is initialized.
9/27/2016 7:55:28 AM	Starting share monitoring
9/27/2016 7:55:32 AM	Changing targetActivationCode health to Ok
9/27/2016 7:55:32 AM	Changing sourceActivationCode health to Ok
9/27/2016 7:55:32 AM	Event log entry written: '6004'.
9/27/2016 7:55:32 AM	Changing to ConnectedState from StoppedState in response to StartSucceededEvent consu...
9/27/2016 7:55:32 AM	Exited StoppedState
9/27/2016 7:55:32 AM	Entered ConnectedState
9/27/2016 7:55:32 AM	Subscribing to engine connection.
9/27/2016 7:55:32 AM	Changing connection health to Ok
9/27/2016 7:55:32 AM	Changing to SynchronizedState from ConnectedState in response to MirrorCompletedEvent...
9/27/2016 7:55:32 AM	Entered ProtectingState
9/27/2016 7:55:32 AM	Entered SynchronizedState
9/27/2016 7:55:32 AM	Persisting shares

The following table identifies the controls and the table columns in the **Job logs** window.

Start

This button starts the addition and scrolling of new messages in the window.

Pause

This button pauses the addition and scrolling of new messages in the window. This is only for the **Job logs** window. The messages are still logged to their respective files on the server.

Copy

This button copies the messages selected in the **Job logs** window to the Windows clipboard.

Clear

This button clears the **Job logs** window. The messages are not cleared from the respective files on the server. If you want to view all of the messages again, close and reopen the **Job logs** window.

Time

This column in the table indicates the date and time when the message was logged.

Description

This column in the table displays the actual message that was logged.

Chapter 7 Simulating protection or recovery

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

Before and after simulating a job, you should gather network and system information specific to Double-Take operations. Use the Double-Take Console to automatically collect this data. It gathers Double-Take log files; Double-Take and system settings; network configuration information such as IP, WINS and DNS addresses; and other data which may be necessary in evaluating Double-Take performance.

1. From the Double-Take Console, on the **Servers** page, right-click the source where you will be running the TDU, select **Gather Support Diagnostics**, and specify a location to store the zipped diagnostics information. It may take several minutes for the diagnostics to finish processing. After it is complete, a .zip file containing the information gathered will be created. The file name is based on the machine name.
2. Establish a protection or recovery job, noting the following caveats.
 - When selecting your target, select the **Diagnostics job** checkbox instead of a target server.
 - When you get to the **Set Options** page in the workflow, some options for your selected job type will not be displayed because they are not applicable. For example, target specific selections will not be displayed because there is no actual target with the TDU.
3. Once you have established your job, you should ideally let it run for several days to gather accurate data for your network and source server usage. The simulation data will be logged to the Double-Take statistics file. See the *Reference Guide* for details on DTStat.
4. After your simulation is complete, repeat step 1 to gather diagnostic again.

Chapter 8 Managing images

Highlight a repository server on the **Servers** page and click **View Server Details** from the toolbar. In the **Tasks** area on the right on the **View Server Details** page, click **Manage images**.

The images stored on a DR target repository server will be listed on the **Images** page. As you highlight an image, any snapshots associated with that image and the virtual disks associated with that image will be listed below.

- **Images, Edit**—Click this button to update the **Image name** and **Image description**.
- **Images, Delete**—You can delete any image that is not part of an active job. If a job is active, delete the job and you will be prompted to keep or delete the image associated with that job. Deleting an image will remove the image from Double-Take recovery processing. Therefore, you cannot select the image for recovery. If you want to be able to recover to an image, do not delete the image.
- **Snapshots, Delete**—You can delete any snapshots from an image for an active or inactive job. Deleting a snapshot will remove the snapshot from Double-Take recovery processing. Therefore, you cannot select the snapshot for recovery. If you want to be able to recover to a snapshot, do not delete the snapshot.

Chapter 9 Special network configurations

Double-Take can be implemented with very little configuration necessary in small or simple networks, but additional configuration may be required in large or complex environments. Because an infinite number of network configurations and environments exist, it is difficult to identify all of the possible configurations. Review the following sections for configuration information for that particular type of network environment.

- *Firewalls* on page 150
- *IP and port forwarding* on page 151

Firewalls

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

- **Double-Take ports**—Ports 6320, 6325, and 6326 are used for Double-Take communications and must be open on your firewall. Open UDP and TCP for both inbound and outbound traffic.
- **Microsoft WMI and RPC ports**—Some features of Double-Take and the Double-Take Console 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.
- **Microsoft File Share and Directory ports**—Double-Take push installations will also rely on File Share and Directory ports, which must be open on your firewall. Check your Microsoft documentation if you need to modify these ports.
 - Microsoft File Share uses ports 135 through 139 for TCP and UDP communications.
 - Microsoft Directory uses port 445 for TCP and UDP communications.

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

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

IP and port forwarding

As outlined in the requirements, Double-Take supports IP and port forwarding in NAT environments with the following caveats.

- Only IPv4 is supported.
- Only standalone servers are supported.
- DNS failover and updates will depend on your configuration
 - Only the source or target can be behind a router, not both.
 - The DNS server must be routable from the target

When setting up a job in an environment with IP or port forwarding, make sure you specify the following configurations.

- Make sure you have added your server to the Double-Take Console using the correct public or private IP address. The name or IP address you use to add a server to the console is dependent on where you are running the console. Specify the private IP address of any servers on the same side of the router as the console. Specify the public IP address of any servers on the other side of the router as the console. This option is on the **Add Servers** page in the **Manual Entry** tab.

The screenshot shows the 'Add Servers' web interface. It has a red header with the title 'Add Servers'. Below the header, there is a sub-header: 'Identify the servers in your environment that you want to manage. The servers you add here appear on the Servers page.' There are two tabs: 'Manual Entry' (selected) and 'Automatic Discovery'. The 'Manual Entry' form contains the following fields: 'Server:' with the value '172.31.206.205', 'User name:' with 'domain administrator', 'Password:' with a masked password, 'Domain:', and 'Management Service port:' with a dropdown set to '1025' and a checkbox for 'Use default port'. An 'Add' button is located at the bottom right of the form. To the right of the form is a table titled 'Servers to be added:' with two columns: 'Server' and 'Details'. The table is currently empty.

- When choosing the target server for your job, you may be prompted for a route from the target to the source. This route, and a port if you are using a non-default port, is used so the target can communicate with the source to build job options. This dialog box will be displayed, only if needed, after you click **Next** on the **Choose Target** page in the job creation wizard.

The screenshot shows a 'Source Address' dialog box with a blue title bar. It contains a yellow warning icon and the following text: 'The Target Server was unable to communicate with the Source Server using any known route. Please provide an alternate IP address that can be used.' Below the text are two input fields: 'Alternate Source IP Address:' with the value '172.31.206.203' and 'Management Service port:' with a dropdown set to '1025' and a checkbox for 'Use default port'. At the bottom are 'OK' and 'Cancel' buttons.

- If you are configuring network monitoring, make sure you specify the port to use for monitoring the Double-Take replication service. This option is in the job creation wizard on the **Set Options** page in the **Failover Monitor** section.

Failover Monitor

Total time to failure: 00:05:00

Consecutive failures: 20

Monitor on this interval: 00:00:10

Network monitoring

Monitor these addresses:

Source IP Address	Replication Service port	
<input type="checkbox"/> 10.10.0.10		
<input checked="" type="checkbox"/> 172.31.216.205 (NAT)	1025 <input type="checkbox"/> Use default port	

Monitoring method:
Network service

Failover trigger:
All monitored IP addresses fail

- When specifying a network route or reverse routes for full server jobs, you can enter a public IP address and then specify ports for the Double-Take Management Service and Double-Take replication service. This option is in the job creation wizard on the **Set Options** page in the **Network Route** section or the **Reverse Protection and Routing** section for full server jobs.

Network Route

Send data to the target server using this route: 172.31.216.205

Management Service port: 1025 Use default port

Replication Service port: 1025 Use default port

Reverse Protection and Routing

Send data to the target server using this route: 172.31.216.205

Management Service port: 1025 Use default port

Replication Service port: 1025 Use default port

Enable reverse protection

A reserved IP address permanently identifies each server so that failover and reverse can both be performed. The reserved IP addresses will not be moved on failover or reverse. These addresses will also be used to route the data in non-NAT environments.

Select a reserved IP address on the source:
10.10.0.10

Send data to source after reverse using this route: 172.31.216.203

Replication Service port: 1025 Use default port

Select a reserved IP address on the target:
20.20.0.10