Double-Take® AVAILABILITY™

Version 7.0
User's Guide

VISION® SOLUTIONS

Notices

Double-Take Availability for Windows User's Guide Version 7.0, Friday, December 20, 2013

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

- **Product Updates**—Check your service agreement to determine which updates and new releases you may be eligible for. Product updates can be obtained from the support web site at http://www.VisionSolutions.com/SupportCentral.
- **Sales**—If you need maintenance renewal, an upgrade activation code, or other sales assistance, contact your reseller/distributor or a Vision Solutions sales representative. Contact information is available on the Vision Solutions Worldwide Locations and Contacts web page at http://www.VisionSolutions.com/Company/Vision-HA-Locations.aspx.
- **Technical Support**—If you need technical assistance, you can contact CustomerCare. All basic configurations outlined in the online documentation will be supported through CustomerCare. Your technical support center is dependent on the reseller or distributor you purchased your product from and is identified on your service agreement. If you do not have access to this agreement, contact CustomerCare and they will direct you to the correct service provider. To contact CustomerCare, you will need your serial number and activation code. Contact information is available on the Vision Solutions CustomerCare web page at http://www.VisionSolutions.com/Support/Support-Overview.aspx .
- **Professional Services**—Assistance and support for advanced configurations may be referred to a Pre-Sales Systems Engineer or to Professional Services. For more information, see the Windows and Linux tab on the Vision Solutions Consulting Services web page at http://www.VisionSolutions.com/Services/Consulting-Services.aspx .
- **Training**—Classroom and computer-based training are available. For more information, see the Double-Take Product Training web page at http://www.VisionSolutions.com/Services/DT-Education.aspx.
- **Documentation**—Please forward any comments or suggestions about this documentation to documentation-Double-Take@VisionSolutions.com.

Contents

Chapter 1 Double-Take Availability overview

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

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

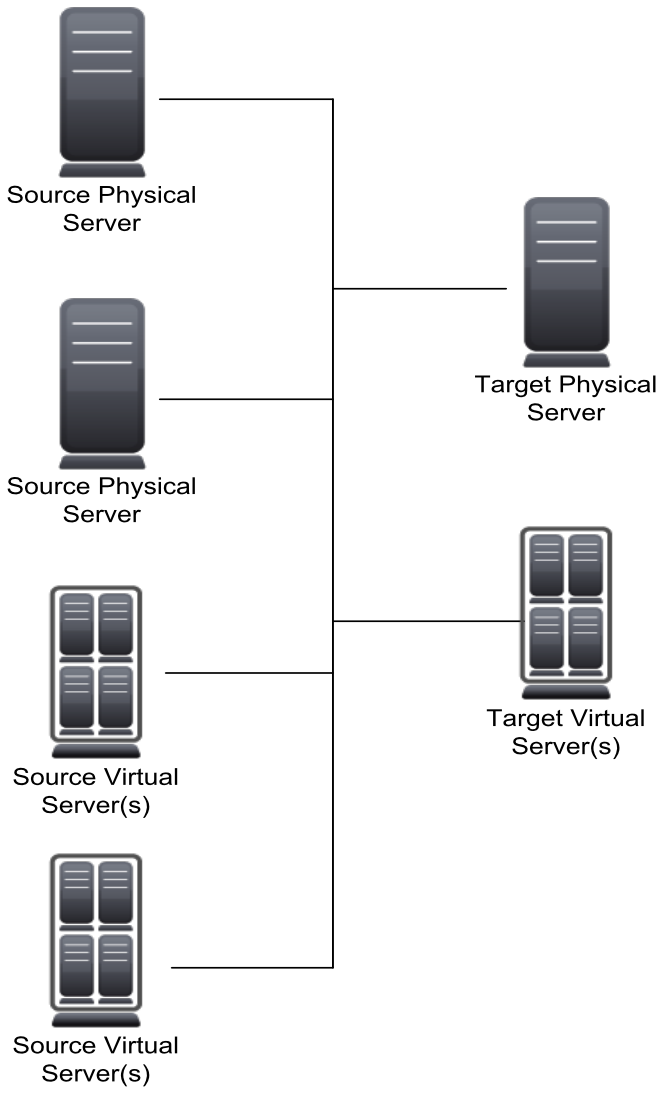

Core operations

Double-Take performs three basic types of operations.

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

Mirroring

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

If subsequent mirroring operations are necessary, Double-Take can mirror specific files or blocks of changed data within files. By mirroring only files that have changed, network administrators can expedite the mirroring of data on the source and target servers. Mirroring has a defined end point when all of the selected files from the source have been transmitted to the target. When a mirror is complete, the target contains a copy of the source files at that point in time.

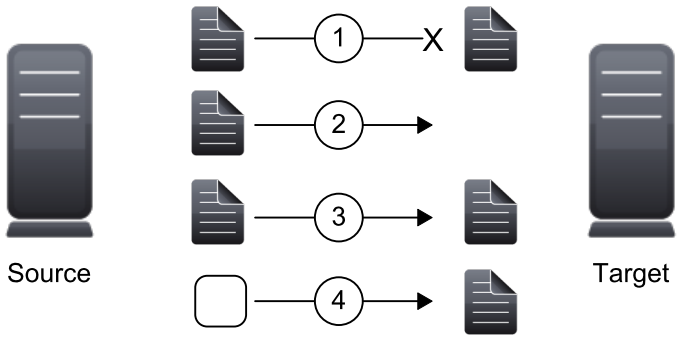

1. Identical files are not mirrored.
2. New files are mirrored.
3. Different files can be mirrored.
4. Checksums can calculate blocks of data to be mirrored.

Replication

Replication is the real-time transmission of file changes. Unlike other related technologies, which are based on a disk driver or a specific application, the Double-Take replication process operates at the file system level and is able to track file changes independently from the file's related application. In terms of network resources and time, replicating changes is a more efficient method of maintaining a real-time copy of data than copying an entire file that has changed.

After a source and target have been connected through Double-Take, file system changes from the user-defined data set are tracked. Double-Take immediately transmits these file changes to the target server. This real-time replication keeps the data on the target up-to-date with the source and provides high availability and disaster recovery with minimal data loss. Unlike mirroring which is complete when all of the files have been transmitted to the target, replication continuously captures the changes as they are written to the source. Replication keeps the target up-to-date and synchronized with the source.

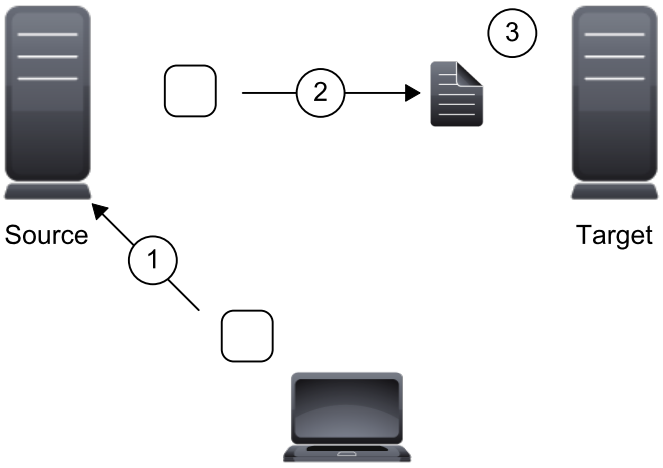

1. A user or application updates part of a file.
2. Only the changed portion of the file is replicated to the target.
3. An up-to-date copy of the file is maintained on the target.

Failover

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

Double-Take monitors the source status by tracking requests and responses exchanged between the source and target. When a monitored source does not respond to the target's requests, Double-Take assumes that the server has failed. Double-Take then prompts the network administrator to initiate failover, or, if configured, it occurs automatically. The failover target assumes the identity of the failed source, and user and application requests destined for the source server or its IP address(es) are routed to the target.

When partnered with the Double-Take data replication capabilities, failover routes user and application requests with minimal disruption and little or no data loss.

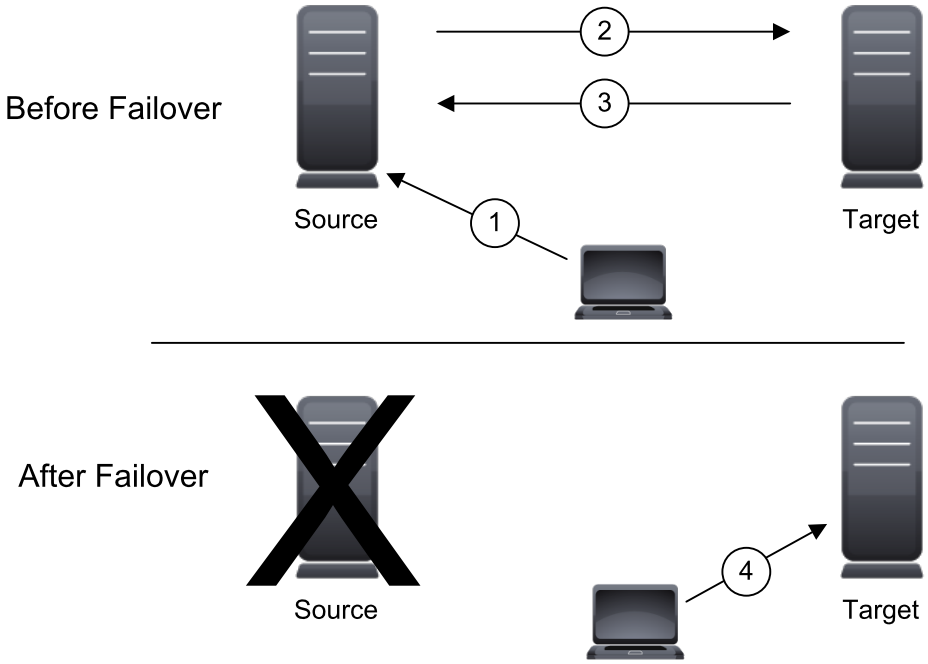

1. User and application requests are sent to the source name or IP address.
2. Data on the source is mirrored and replicated to the target.
3. The target monitors the source for failure.
4. In the event the source fails, the target stands in for the source. User and application requests are still sent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

Double-Take Availability workloads

In addition to selecting your own files and folders that you want to protect, Double-Take can protect specific types of workloads to meet your protection and business goals.

Full server protection

Full server protection provides high availability for an entire server, including the system state, which is the server's configured operating system and applications. You identify your source, which is the server you want to protect, and your target, which is the server that will stand-in for the source in the event the source fails. Double-Take monitors the source for a failure, and if it fails, the target will stand-in for the source by rebooting and applying the source system state on the target. After the reboot, the target becomes the source.

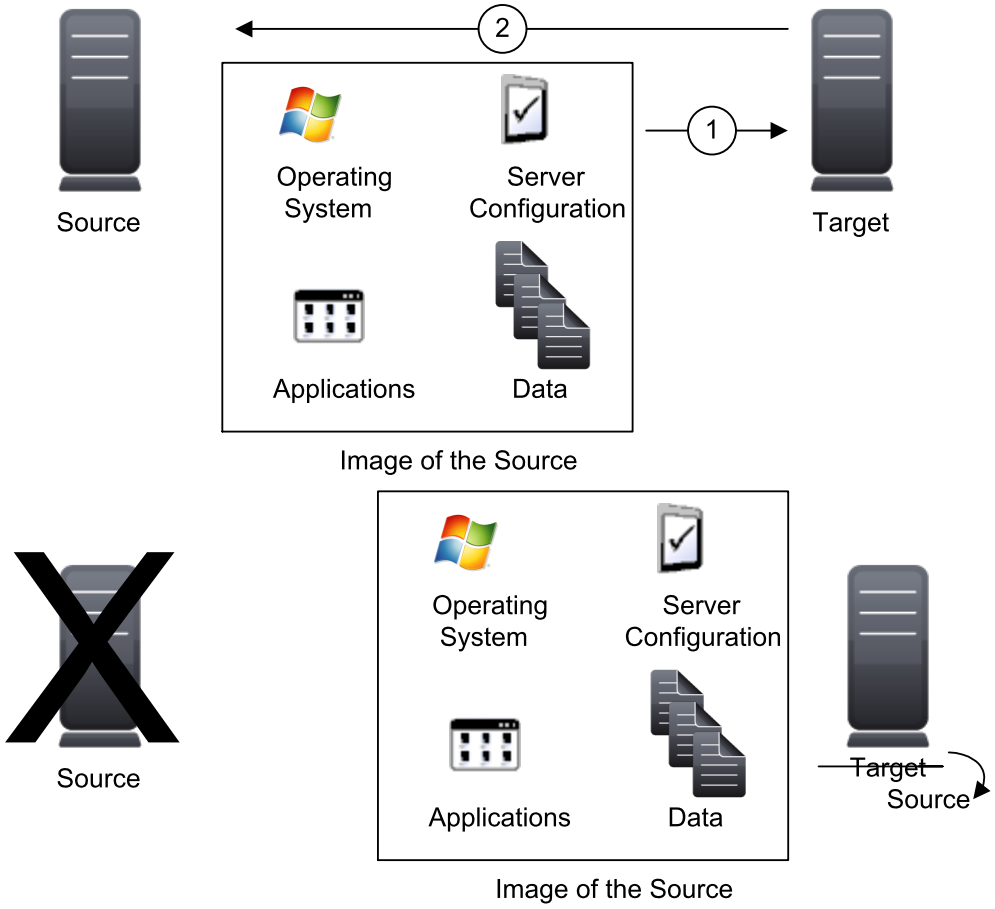

1. The source data and system data, together a total image of the source, are mirrored and replicated to the target.
2. The target monitors the source for failure.
3. In the even the source fails, the source's system state is applied when the target is rebooted. After the reboot, the target is now the source, in both identity and with the source data.

Application protection

Application protection provides high availability for Microsoft Exchange or Microsoft SQL Server. You identify your source, which is the server running the application, and your target, which is the server that will stand-in for the source in the event the source fails. Double-Take will gather information from your environment (application configuration, Active Directory, DNS, and so on) about the application being protected and automatically protect the application. Double-Take monitors the source server or the application services for a failure. If it fails, the target will stand-in for the source. End-users continue to access the application, now running on the target.

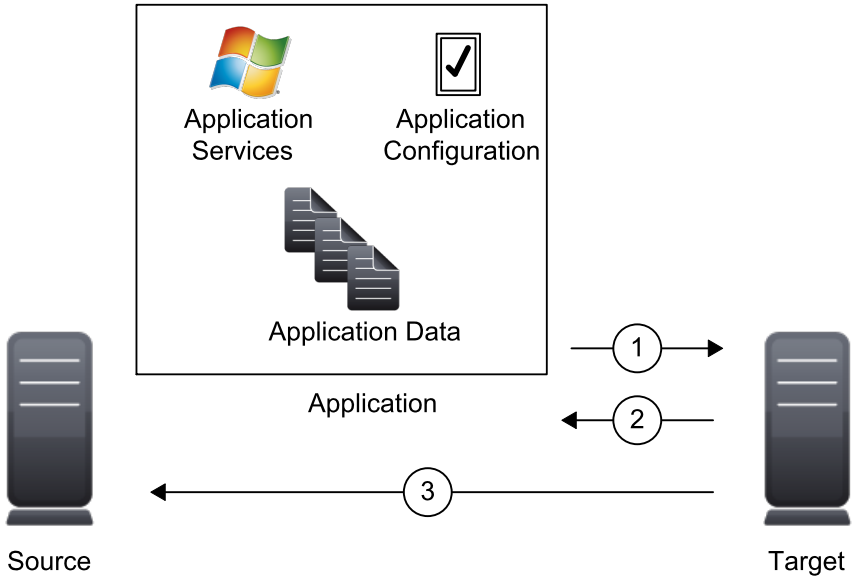

1. The configuration is sent to the target and then application data is mirrored and replicated to the target.
2. The target can monitor the application for failure.
3. The target can monitor the source for failure.

Virtual protection

Virtual protection provides high availability to Hyper-V or ESX virtual servers. You identify your source, which is the server you want to protect. Your source can be a physical server, a virtual machine where you want to protect the data within the guest operating system, or a virtual machine where you want to protect the host-level virtual disk files (.vhd or .vhdx files). Your target is a Hyper-V or ESX server that will host a virtual machine that is a replica of the source. Double-Take monitors the source for a failure. In the event of a source failure, the replica virtual machine on the target can stand-in allowing end-users to continue accessing data and/or applications.

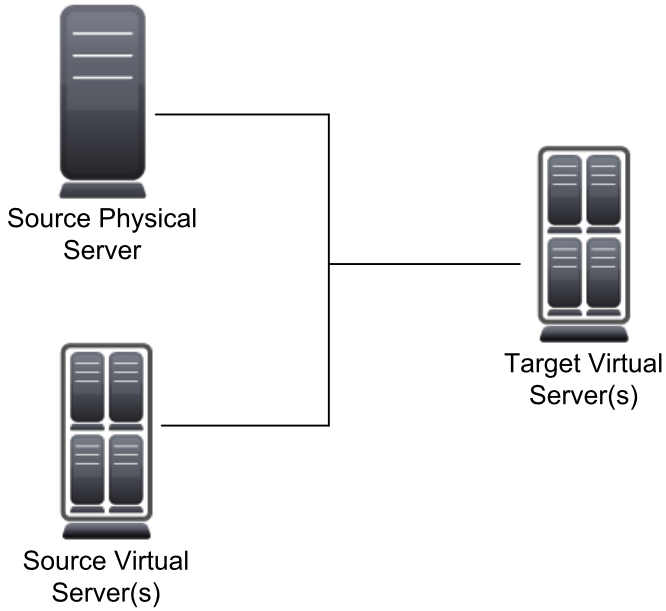

You can protect virtual workloads in many different configurations. See *Selecting a protection type* on page 165 for a decision tree to help you determine the best virtual protection for your environment.

GeoCluster protection

In a standard cluster configuration, a single copy of data resides on a SCSI disk that is shared between cluster nodes. Data is available without users knowing which node owns a cluster resource. MSCS handles failover between nodes of the cluster.

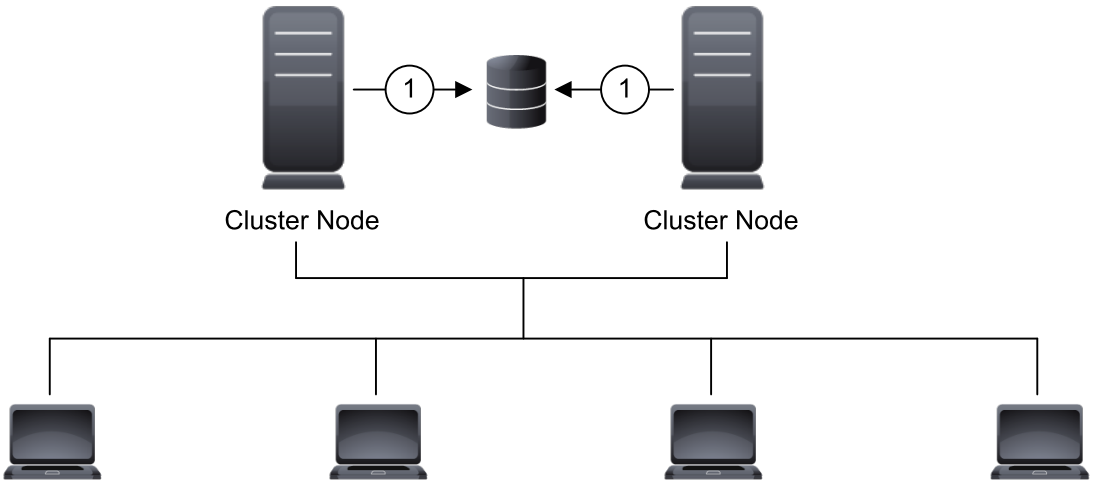

1. The source cluster nodes share data from a single SCSI disk.

In a GeoCluster configuration, data is stored on volumes local to each node and replicated to each node in the cluster using Double-Take. This eliminates the single point of failure of a standard cluster (shared disk) configuration. With GeoCluster, resources and groups are handled in the same manner as a standard cluster. Instead of assigning one group by SCSI drive, you assign one group per logical volume. If a server, disk, group, or network interface should fail, MSCS relocates the failed group to another node, which contains the replicated copy of the data, thus maintaining availability.

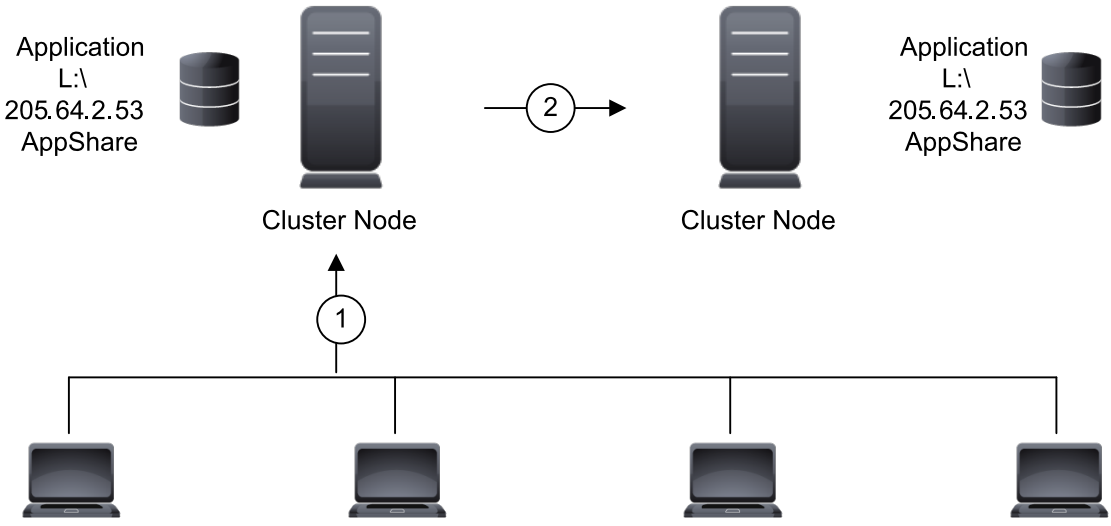

1. Users access data from the owning node.
2. Data is mirrored and replicated between nodes of the cluster.

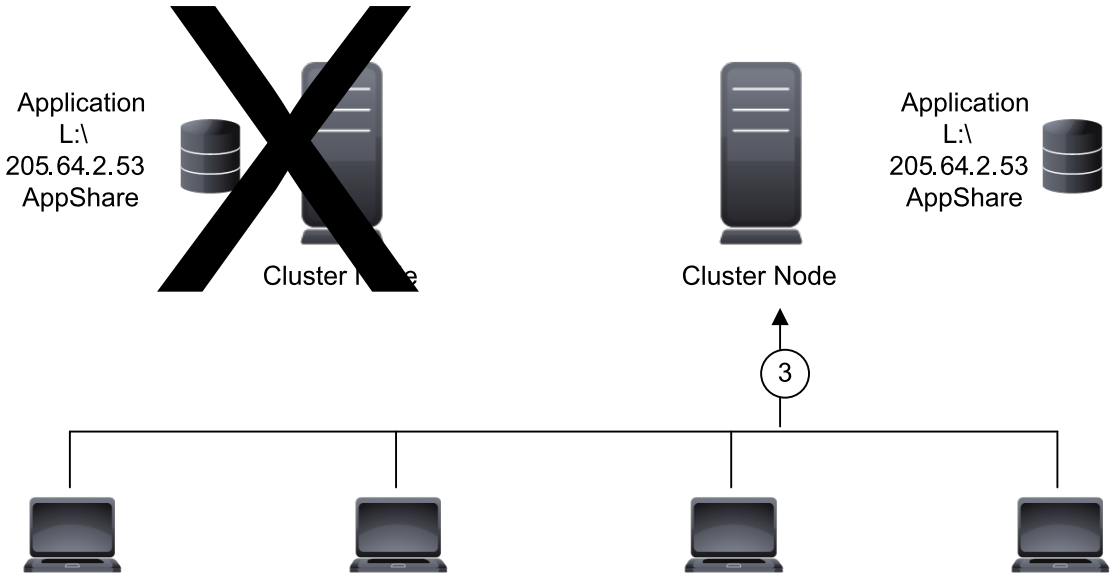

3. In the event the owning node changes, users access data from the new owning node.

Supported configurations

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

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

- See *One to one, active/standby* on page 17
- See *One to one, active/active* on page 18
- See *Many to one* on page 19
- See *One to many* on page 20
- See *Chained* on page 21
- See *Single server* on page 22

One to one, active/standby

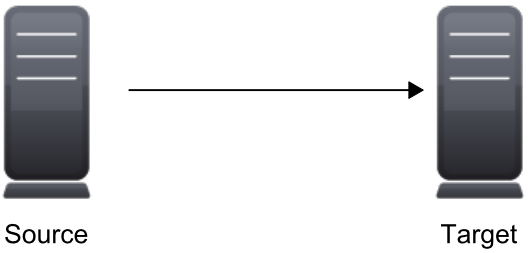

Description

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

Applications

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

Considerations

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

One to one, active/active

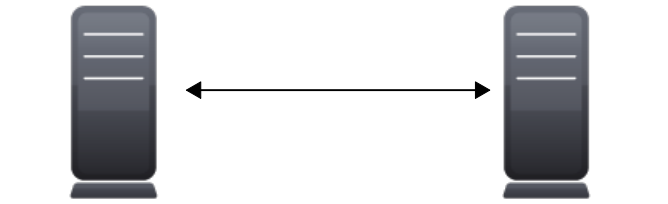

Description

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

Applications

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

Considerations

- Coordination of the configuration of Double-Take and other applications can be more complex than the one to one active/standby configuration.
- During replication, each server must continue to process its normal workload.
- Administrators must avoid selecting a target destination path that is included in the source's protected data set. Any overlap will cause an infinite loop.
- To support the production activities of both servers during failover without reducing performance, each server should have sufficient disk space and processing resources.
- Failover and failback scripts must be implemented to avoid conflict with the existing production applications.
- You must pause the target when backing up database files on the target.

Many to one

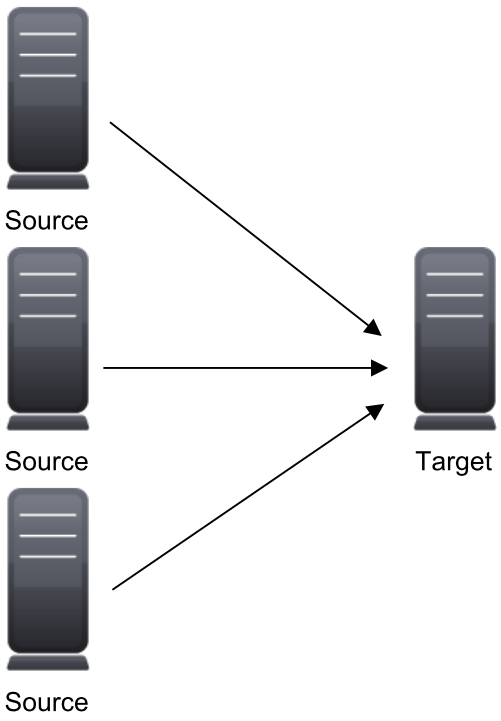

Description

Many source servers are protected by one target server.

Applications

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

Considerations

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

One to many

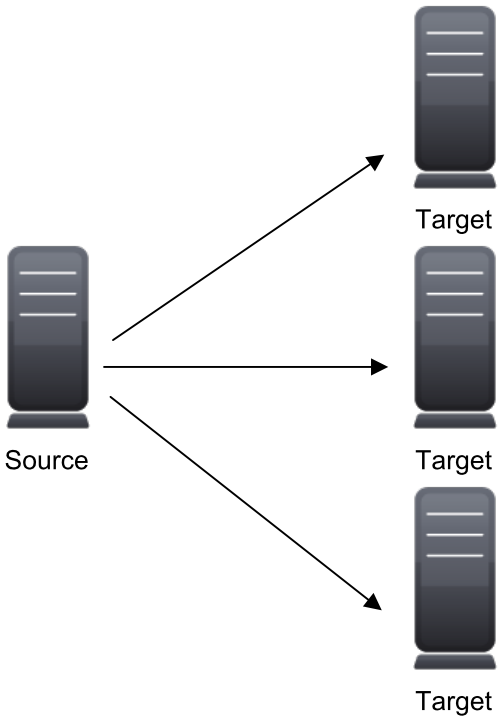

Description

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

Applications

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

Considerations

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

Chained

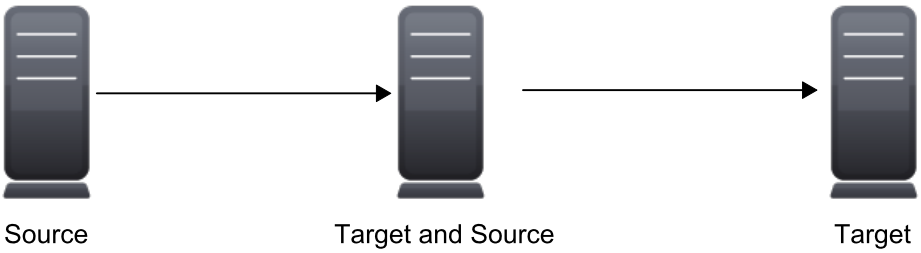

Description

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

Applications

This is a convenient approach for integrating local high availability with offsite disaster recovery. This configuration moves the processing burden of WAN communications from the source server to the target/source server. After failover in a one to one, many to one, or one to many configuration, the data on the target is no longer protected. This configuration allows failover from the first source to the middle machine, with the third machine still protecting the data.

Considerations

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

Single server

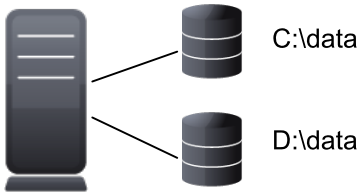

Source and Target

Description

Source and target components are loaded on the same server allowing data to be replicated from one location to another on the same volume or to a separate volume on the same server. These could be locally attached SCSI drives or Fibre Channel based SAN devices.

Applications

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

Considerations

- This configuration does not provide high availability for the entire server.
- This configuration must be configured carefully so that an infinite loop is not created.

Chapter 2 Core Double-Take requirements

Each Windows server must meet the minimum core requirements below. Your servers may need to meet additional requirements depending on the job type will be using. See the requirements section for each of the job types for those specific requirements.

- **Operating system**—Double-Take for Windows supports Windows 2003, 2008 and 2012, with the following caveats.
  - Windows 2003 operating systems require Service Pack 1 or later.
  - Server Core 2008 is not supported for any job type.
  - Server Core 2008 R2, 2012, and 2012 R2 are supported for files and folders, full server, full server to Hyper-V, V to Hyper-V, and agentless Hyper-V jobs but with caveats. See the requirements sections for those job types for Server Core details.
  - Hyper-V is supported for full server to Hyper-V, V to Hyper-V, and agentless Hyper-V jobs. See the requirements sections for those job types for detailed Hyper-V requirements.
  - There are multiple Double-Take editions that correspond to the various Windows operating system editions. Your reseller or distributor can assist you in selecting the right Double-Take edition for your operating system and protection needs.
  - If you are using Windows 2008 Hyper-V R2 with Cluster Shared Volumes (CSV), you can use Double-Take within the guest. However, you will be unable to protect the virtual machines at the host level using Double-Take. Although CSV is not supported for host-level protection for Windows 2008, you can use host-level protection for non-CSV virtual machines on a CSV configured cluster. See *Agentless Hyper-V requirements* on page 604 for more information.
  - Cluster Shared Volumes (CSV) support at the host-level is for Windows 2012 only.
  - Even though the Double-Take for Windows Foundation Edition is for Windows Foundation Server, Small Business Server, Storage Server, and Essential Business Server editions, you can install the Double-Take for Windows Foundation Edition on a server running higher Windows operating systems, so you do not have to pay extra for an upgraded Double-Take license. When the Foundation Edition is installed on a server running higher Windows operating systems, the following limitations will apply. 1) The server will function in a target role only. 2) The target-only server can only protect a source that is running the Double-Take for Windows Foundation Edition and one of the operating system editions listed above. 3) Full server protection is the only supported method of failover.
- **File system**—Double-Take supports the NTFS and ReFS file system. FAT and FAT32 are not supported. For detailed information on other file system capabilities, see *Mirroring and replication capabilities* on page 28.
- **System memory**—The minimum system memory on each server should be 1 GB. The recommended amount for each server is 2 GB.
- **Disk space for program files**—This is the amount of disk space needed for the Double-Take program files. The amount depends on your operating system version and your architecture (32-bit or 64-bit) 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 and appliances 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). See *Deactivating licenses* on page 58. 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 *Activating a single license* on page 54 for details on completing a host-transfer.

- **Time**—The clock on the source and target servers must be within a few minutes of each other, relative to UTC. Large time skews 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. (Some job types allow you to add DHCP addresses for failover monitoring, although only after a job has already been created. Keep in mind that depending on your failover configuration, a source reboot may or may not cause a failover but having a new address assigned by DHCP may also cause a failover.)
  - 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 later add IPv6, you will need to manually modify the DefaultProtocol server setting. See *Server and job settings* on page 104 for details.
  - For some job types, IPv4 is the only supported version. See the requirements for each job type.
  - IPv6 is only supported for Windows 2008 and 2012 servers.
  - 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.
- **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**—Some job types allow you to update DNS at failover time. To be able to use DNS

updates, your environment must meet the following requirements.

- The source and target servers must be in the same domain.
- At failover time, the target must be able to reach the DNS servers that you want to update.
- For workgroup environments, both the source and target server must be in the workgroup.
- Each server's network adapter must have the DNS suffix defined, and the primary DNS suffix must be the same on the source and target. 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.

DNS updates are not supported for Server Core servers or NAT environments.

- **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 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 811 for instructions on handling firewalls in your environment.
- **Windows Management Instrumentation (WMI)**—Double-Take is dependent on the WMI service. If you do not use this service in your environment, contact technical support.
- **Snapshots**—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 operating system**—Your servers must be running, at a minimum, Windows 2003 Service Pack 1. You should upgrade to Service Pack 2 or later so that several Microsoft patches that address memory leaks in the Volume Shadow Copy service are applied. If you do not have Service Pack 2 installed, you will need to review the patches available on the Microsoft web site and install those that correct the Volume Shadow Copy service memory leaks.
  - **Snapshot data**—Snapshots are taken at the volume level. For example, if your job is protecting D:\data and E:\files, the snapshot will contain all of the data on both the D: and E: volumes. If your job is only protecting D:\data (E:\files exists but is not included in the job), the snapshot will only contain the D: volume.
  - **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.
  - **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.

Additionally, Volume Shadow Copy will not revert snapshots of a volume with operating system files, therefore Double-Take is also unable to revert a volume with operating system files.

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.

- **Other snapshot functionality**—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.

- **Clusters**—If your job type supports clustering, make sure your cluster meets the following general requirements in addition to any job specific cluster requirements.
  - **Best practices**—You should carefully review Microsoft documentation and resources for properly configuring your cluster before implementing Double-Take on a cluster. The Microsoft TechNet articles Failover Clusters and Installing and Upgrading Cluster Nodes are two resources you can start with. There are many other resources available on the Microsoft TechNet web site.
  - **Networking**—The following networking requirements apply to your cluster.
    - You must have TCP/IP connections between nodes.
    - Multiple networks are recommended to isolate public and private traffic.
    - The private network should be a unique subnet so that Double-Take will not attempt to use an unreachable private network.
    - Your network can contain direct LAN connections or VLAN technology.
    - For Windows 2003, the cluster nodes must be on the same logical IP subnet.
    - For Windows 2003, the maximum round trip latency between nodes should be no more than ½ second.
  - **Domain**—The cluster nodes must be members of the same domain.

- **DNS**—Forward and reverse lookups must be implemented on the primary DNS server for the cluster name and individual nodes.
- **Cluster service account**—For Windows 2003 clusters, use the same cluster service account on source and target clusters.
- **Double-Take disk queue**—Ensure that the disk queue is not on a Physical Disk resource.
- **Volumes**—The source and target should have identical drive mappings.
- **Licensing**—Each node in the cluster must have a valid Double-Take Availability activation code.
- **Resource registration**—In some cases, the Double-Take cluster resources may not be registered automatically when Double-Take is installed. You can manually register the resources by running DTResUtility.exe, which is installed in the \Windows\Cluster directory.
- **Third-party storage**—Third-party storage resources are not supported.

Mirroring and replication capabilities

For Windows source servers, Double-Take mirrors and replicates file and directory data stored on any NTFS or ReFS 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. You must mirror and replicate to like file systems. For example, you cannot use NTFS to ReFS or ReFS to NTFS. You must use NTFS to NTFS or ReFS to ReFS. Additionally, you cannot have ReFS volumes mounted to mount points in NTFS volumes or NTFS volumes mounted to mount points in ReFS volumes.
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 checksum option when performing a difference mirror or verification to ensure that all alternate data streams are compared correctly.
    d. If your alternate streams are read-only, the times may be flagged as different if you are creating a verification report only. Initiating a remirror with the verification will correct this issue.
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 failover occurs any pending transactions on the target will be aborted.

   d. Double-Take restore functions as a mirror, except the roles of the source and target are reversed. The transacted view of the data on the target is restored to the source. As

pending transactions are committed or aborted on the target, Double-Take restores any necessary changes to the source. Once all pending transactions are completed, the restoration is complete and replication will continue from the target to the source.

e. If you have restored your data before starting the failback process, make sure the restoration process does not have pending transactions and is complete before starting failback. If you are restoring your data after the failback the process has completed, users will not be accessing the data once failback occurs, so there are no opportunities for pending transactions.

10. Double-Take supports Windows 2008 and 2012 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 (set to 1) on the source but disabled (set to 0) 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. Note that the Windows 2012 default value for this setting is disabled (set to 0).
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, failover 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 are running Symantec version 10 or later, you may receive Event message 16395 indicating that Double-Take has detected a hard link. Symantec uses a hard link to recover from a virus or spyware attack. Double-Take does not support hard links, therefore, the Event message is generated, but can be disregarded.
19. 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 sparse on the target and has all the data).

20. 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 failover in order to access the data. Additionally, you cannot have reparse points in your data set if you are using same server protection because the server is functioning as both a source and target.
21. If you are using an archiving solution, do not archive any files after failover. Archiving files after failover could cause corruption.
22. 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, remirror, or a restoration mirror) is running. Windows updates that occur during a mirror may cause data integrity issues on the target.
    b. You must resolve any Windows Update incomplete operations or errors before failover or failback. (Check the windowsupdate.log file.) Also, do not failover or failback if the target is waiting on a Windows Update reboot. If failover 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.
23. 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 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.
    c. After failover, 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.
24. If you are using Windows storage pools on your source, you must create the storage pool on the target before failover.

Chapter 3 Installation and activation

The installation and activation for Double-Take servers is a two-part process. First, you will have to install the software on your servers. Second, you will have to activate your licenses. Choose how you want to install and activate your servers using either of the following methods.

- **Single server installation and activation**—You can install and activate your servers one at a time. In this scenario, you will install Double-Take on one server and then you will activate the license on that server. You can repeat this process for other servers.
- **Multiple server installation and activation**—You can install and activate multiple servers at one time. In this scenario, you will need install Double-Take on all of your servers and then you will activate the licenses for all of the servers at once.

The term installation is being used generically to represent a new Double-Take installation or a Double-Take upgrade.

If you have received an evaluation activation code in order to test Double-Take before purchasing a license, you will not need to perform the activation steps. Your evaluation activation code will automatically disable all Double-Take functionality when the evaluation period expires. You will then need to update to a valid license and go through the activation process.

If you are going to be creating a V to Hyper-V or V to ESX job (see *Selecting a protection type* on page 165 for details on which job to create), the job creation process can automatically push the Double-Take installation out to the servers, if you have the Double-Take Availability Virtual Guest Edition license. After the job is created, you will have 14 days to activate the licenses using either the single server method or the multiple server method. If you do not have a Double-Take Availability Virtual Guest Edition license, you will need to perform the installation outside of the job creation process using the single server or multiple servers method.

If you are protecting a GeoCluster configuration, where data is stored on volumes local to each node and replicated to each node in the cluster, you should confirm or configure your cluster configuration before beginning the installation. See *Configuring a cluster for GeoCluster* on page 656.

See *Installation* on page 34 and *License management and activation* on page 48 for details on the various installation and license and activation methods.

Installation

Review the *Installation notes* on page 35 before beginning your installation. Then choose from one of the following installation methods.

- *Installing using the installation wizard* on page 37—Use these instructions to install on a single Windows server using the installation wizard. The wizard will guide you step by step through the installation process.
- *Installing using the command line utility* on page 40—Use these instructions to install on a single or multiple Windows server using the command line utility. This utility automates the installation process by running an unattended, or silent, installation. It allows you to pass parameters through to the installation program instead of entering information manually through the wizard.
- *Installing using the Double-Take Console* on page 44—Once you have the Double-Take Console installed on a Windows machine, you can use it to push the installation out to your other servers. The push installation is a full, client and server installation on Windows servers and a full server installation on Linux servers.

Installation notes

Review the installation notes below before beginning an installation or upgrade.

- Because Double-Take has operating system dependent files, if you are upgrading your operating system (to a new major version, not a service pack) and have Double-Take installed, you must remove Double-Take prior to the operating system upgrade. Uninstall Double-Take, perform the operating system upgrade, and then reinstall Double-Take.
- During the installation, DTInfo is installed. This is a utility that can be run to collect 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.
    - On Windows servers, the DTInfo utility is DTInfo.exe, and it is installed to the Double-Take installation directory. After running the executable, a zip file is automatically created with the information gathered. You can also collect this information from the Double-Take Console, storing the resulting zip file on the console machine . See *Managing servers* on page 65.
    - On Linux source servers, the DTInfo utility is DTInfo.sh, and it can be run from DTSetup. After running the script, a .tar.gz is automatically created with the information gathered. You must have root (or uid 0 equivalent) to execute the diagnostics or to copy or read the resulting file.You can also collect this information from the Double-Take Console, storing the resulting file on the console machine . See *Managing servers* on page 65.
    - On the Double-Take Linux virtual recovery appliance, DTInfo can be run manually from /usr/bin. It should be executed as DTInfo.sh -f. After running the script, a .tar.gz is automatically created with the information gathered. You can also collect this information from the Double-Take Console, storing the resulting file on the console machine . See *Managing servers* on page 65.
- The following notes are specific to Windows servers.
    - Since Double-Take installs device drivers, it is recommended that you update your Windows Recovery Disk, before installing or making changes to your servers. For detailed instructions on creating a recovery disk, see your Windows reference manuals. Make sure that you select the option to back up the registry when building the repair disks.
    - If you are installing to a drive other than the drive which contains your system TEMP directory, the Microsoft Windows Installer will still load approximately 100 MB of data to the TEMP directory during the installation. If you do not have enough disk space on the drive that contains the TEMP directory, you may need to change where it is located.
    - If during the installation you receive the message that the wizard was interrupted before the installation could be completed, you will need to delete the registry value DefaultAccessPermissions under the HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Ole key in order to install Double-Take. This registry setting denies permissions for retrieving and setting property values. Deleting this registry setting will not have a negative impact on your server.
- Double-Take 7.0 is interoperable back to version 5.3 for Windows servers and 6.0 for Linux servers but is restricted to the following limitations. The Double-Take clients can only control the same or older releases. To accommodate rolling upgrades, older sources can connect to newer targets, but newer sources cannot connect to older targets.
    - **5.3 console**—Supports 5.3 source and target, but does not support 6.0 or 7.0 source or target

- **6.0 console**—Supports 5.3 or 6.0 source and target as long as the target is the same or newer than the source, but does not support 7.0 source or target
- **7.0 console**—Supports 5.3, 6.0, or 7.0 source and target as long as the target is the same or newer than the source

Check the requirements for each individual job type for specific upgrade and console configurations.

- When performing a rolling upgrade, update the target servers first. After the upgrade is complete, the sources will automatically reconnect to the targets. Upgrade the sources when convenient.
- If you are using a chained configuration, update the last target first, then update the middle server acting as both a source and target, and update the production source last.
- If you are using a configuration where the source is an older version than the target, you will not be able to restore from the newer version target back to the older version source. You must upgrade the source to the same version as the target before restoring.
- Use the following procedure to upgrade Double-Take on a Windows multi-node cluster. If both your source and target are clusters, use the following procedure on the target cluster first, then on the source. If you are protecting Hyper-V virtual machines at the host-level using a cluster configuration, you cannot upgrade. You must delete the existing job, upgrade all of your nodes, and then re-create the job.
  1. Move all cluster resources to one node.
  2. Upgrade to the new version of Double-Take on all of the other nodes.
  3. Move the cluster resources to one of the upgraded nodes.
  4. Upgrade to the new version of Double-Take on the last node.
  5. If needed, move the cluster resources to the desired nodes.
- Use the following procedure to upgrade Double-Take on a Windows single-node cluster. If both your source and target are clusters, use the following procedure on the target cluster first, then on the source.
  1. Take your cluster resources offline.
  2. Upgrade to the new version of Double-Take.
  3. After the upgrade is complete, bring the resources back online.
- If you are upgrading and are currently using the GeoCluster Replicated Disk as the quorum resource, you will need to select another quorum resource before upgrading. See *Configuring a cluster for GeoCluster* on page 656 for more information.
- During an installation or upgrade, if the GeoCluster Replicated Disk resource files fail to register with the cluster, use the DTResUtility, located in the \windows\cluster directory, to manually register the resources.
- If you have upgraded the guest Windows operating system on a Hyper-V virtual machine that you intend to protect at the guest level, you must also upgrade Integration Services on that virtual machine.

Installing using the installation wizard

Make sure you have reviewed the *Installation notes* on page 35, and then use these instructions to install Double-Take Availability or upgrade an existing Double-Take Availability installation.

1. Close any open applications.
2. Start the installation program using the appropriate instructions, depending on your media source.
   - **Physical media**—If auto-run is enabled, the installation program will start automatically. To manually start the program, run autorun.exe from your physical media.
   - **Web download**—Launch the .exe file that you downloaded from the web.

   

   If you are installing on Server Core, copy the physical media files or web download file to the Server Core machine using a UNC share, and then launch the installation program from the Server Core machine. The installation screens will display on the Server Core machine.

3. When the installation program begins, the Autorun appears allowing you to install software and view documentation and online resources. To install Double-Take Availability, select the **Install Double-Take Availability** link.
4. Depending on your version of Windows and the components you have installed, you may be prompted to install one or more Visual C++ security updates. These components are required. If you do not see this screen, you server already has the security updates.
5. Also depending on your version of Windows and the components you have installed, you may need to install or enable Microsoft .NET Framework. If you do not see this screen, your server already has the appropriate version of Microsoft .NET. You must install or enable Microsoft .NET before installing Double-Take. Select **Yes** to install or enable Microsoft .NET and click **Continue**.
6. If you are upgrading from a previous version of Double-Take RecoverNow, and you have any Double-Take archived or deduplicated files, you will be prompted to recall and restore those files manually. Double-Take RecoverNow will automatically be uninstalled, however you must recall and restore those files first.
7. You will be given the opportunity to check for a more recent version of the software.
   - If you do not want to check for a later version, select **No** and click **Next**.
   - If you want to check for a later version, select **Yes** and click **Next**. The installation program will establish an Internet connection from your server to the Vision Solutions web site.
     - If later versions are found, they will be listed. Highlight the version you want and either download that version and install it automatically or download that version and exit the installation. (If you exit the installation, you can run the updated installation later directly from the location where you saved it.)
     - If no later versions are found, continue with the current installation.
     - If an Internet connection cannot be established, continue with the current installation or install a previously downloaded version.
8. If you are upgrading, review the upgrade note.
   - Any jobs that were created in legacy Double-Take consoles, including Replication Console, Full Server Failover Manager, Application Manager, Double-Take Move Console, or

DTCL, will no longer function once the upgrade is complete.

- Any jobs (except full server to ESX, full server to Hyper-V, V to ESX, or V to Hyper-V) created in the version 5.3 Double-Take Console will be upgraded and will continue to function normally.
- If you are protecting Hyper-V virtual machines at the host-level using a cluster configuration, you cannot upgrade. You must delete the existing job, upgrade all of your nodes, and then re-create the job.

9. Click **Next** to continue.
10. If you are upgrading from Double-Take RecoverNow with TimeData, select if you want to uninstall the TimeData components. You can also remove Ontrack PowerControls if you have that installed. The uninstall may be time consuming while it removes the associated SQL instance and deletes the continuous data protection storage bins. Additionally, the TimeData uninstall may require a reboot. If you choose not to uninstall TimeData or Ontrack PowerControls, Double-Take RecoverNow will still be uninstalled, and these programs will no longer function after the upgrade. If you do not remove the products, you will have to manually remove them.
11. Review the Vision Solutions license agreement. You must scroll through and review the entire license agreement. You must accept the license agreement in order to continue with the installation program. Click **Next** to continue.
12. Review the activation notice. Most Double-Take licenses require activation after installation for full product functionality. Failure to activate licenses that require it will cause your Double-Take jobs to fail. See *License management and activation* on page 48 for more details.
13. Click **OK** to continue.
14. Select the type of installation you would like to perform on this machine.

    - **Client and Server Components**—This option installs both the client and server components. The server components are required for systems that will function as a source or target. The server requires an activation code for the service to run. The client does not require an activation code, but it is required to administer this and other Double-Take servers throughout the organization.
    - **Client Components Only**—This option installs only the client components. The client components do not require an activation code, but are required to administer Double-Take servers throughout the organization.
    - **Server Components Only**—This option installs only the server components. The server components are required for systems that will function as a source or target. The server requires an activation code for the service to run.

    

    If you are installing on Server Core, you will only be able to select the **Server Components Only** installation. You will not be able to run the client components from the Server Core machine. The client will have to be run from another machine.

15. If desired, specify where the Double-Take files will be installed by clicking **Change**, specifying a location, and then clicking **OK**.
16. Click **Next** to continue.
17. If the machine where you are installing has Windows Firewall enabled, you will be given the opportunity to open and reassign any firewall ports for Double-Take use.

- **Open only the ports that are not in use**—This option will open any firewall ports that are not in use. The ports that are opened will be assigned to Double-Take.
- **Open all ports, reassigning the ports in use to Double-Take**—This option will open all necessary firewall ports, reassigning any to Double-Take as needed.
- **Do not configure the Windows Firewall at this time**—This option will not modify any firewall ports. If you select this option, you will have manually modify your firewall settings for Double-Take processing.

18. Click **Next** to continue.

If you selected a client only installation, continue with step 24.

---

19. You will be prompted to enter your activation code information. Your **Activation Code** is a 24-character, alpha-numeric activation code which applies the appropriate license to your installation. Enter your code and click **Add**.

If you are installing Double-Take for the first time, you will not be required to add an activation code, however the product will not work without one. If you are upgrading an existing version, you must have an activation code to continue the installation.

---

20. Click **Next** to continue.
21. Double-Take uses system memory for Double-Take processing. The maximum amount is dependent on the server hardware and operating system. **Specify the Amount of system memory to use**, which is the maximum amount of system memory that Double-Take can use.
22. When the allocated Double-Take system memory is exhausted, Double-Take will queue to disk. If you want to disable disk queuing, select **Do not use disk queue**. Ideally, you should use disk queuing. Specify the **Queue folder**, which is the location of the disk queue. By default, the size of the disk queue is set to **Unlimited disk queue**, which will allow the queue usage to automatically expand whenever the available disk space expands. If desired, you can select **Limit disk space for queue** and specify a fixed disk space amount. You can also specify the **Minimum free disk space**, which 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**. (See *Double-Take queue* on page 88 for additional guidelines on selecting appropriate queue settings.)
23. Click **Next** to continue.
24. The Double-Take security information screen appears next. Review this information and click **Next** to continue with the installation.
25. If you are satisfied with the selections you have made and are ready to begin copying the Double-Take files, click **Install**.
26. After the files have completed copying, click **Finish** to exit the installation program.

Installing using the command line utility

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

The automatic installation is only available for new installations. It does not support upgrades.

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

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

DTSetupType

- **DTNT**—Both the Double-Take server and client components will be installed.
- **DTCO**—Only the Double-Take client components will be installed.
- **DTSO**—Only the Double-Take server components will be installed.

If you are installing on Server Core or Windows Hyper-V Server (standalone), the setup type will be server components only regardless of your setting.

DTActivationCode

A 24 character, alpha-numeric activation code which applies the appropriate license to the server. Multiple activation codes can be separated by a semi-colon.

DoubleTakeFolder

Any valid path specifying the location of the Double-Take files

QMemoryBufferMax

Any integer representing the amount of system memory, in MB, that Double-Take can use

DiskQueueFolder

Any valid path to the location of the disk-based queue

DiskQueueMaxSize

Any integer representing the amount of disk space, in MB, to use for disk-based queuing or the keyword **UNLIMITED** which will allow the queue usage to automatically expand whenever the available disk space expands

DiskFreeSpaceMin

Any integer representing the amount of disk space, in MB, that must remain free at all times

DTServiceStartup

- **Y** or **1**—Start the Double-Take service automatically
- **N** or **0**—Do not start the Double-Take service automatically

This parameter is not applied if your **DTSetupType** is DTCO.

Port

Any integer between 1024 and 65535 that identifies the primary port used for Double-Take communications

Set_FWPort

- **Y** or **1**—Set the Double-Take Windows firewall port exclusions
- **N** or **0**—Do not set the Double-Take Windows firewall port exclusions

You must have Microsoft .NET installed or enabled (depending on your operating system) on the server before starting the automatic installation.

If you are using Windows 2008 or 2012, but you are not using the built-in administrator account, User Access Control will prompt you to confirm you want to install Double-Take. To work around this issue, use the built-in administrator account when you are installing to each server. You may also disable User Access Control, if that is acceptable for your environment.

Installing or upgrading automatically to a local machine

1. Create a temporary installation directory on the server. For example, create c:\temp_install.
2. Use the following steps if you downloaded your software from the web.
    a. Unzip the .exe file that you downloaded to another temporary directory.
    b. Locate the subdirectory under \setup\dt that is appropriate for your architecture, either i386 or x64.
    c. Copy the files from the \setup\dt\i386 or \setup\dt\x64 directory to your temporary installation directory.
3. Use the following steps if you have a DVD.
    a. Locate the subdirectory under \setup\dt that is appropriate for your architecture, either i386 or x64.
    b. Copy the files from the \setup\dt\i386 or \setup\dt\x64 directory to your temporary installation directory.
4. Remove the read-only attributes from the files in the temporary installation directory.
5. Make a backup copy of the default DTSetup.ini file in the temporary installation directory.
6. Edit DTSetup.ini as needed using the values described in the previous table.
7. Run the following case-sensitive command from the temporary installation directory.

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

The command must be run from the temporary installation directory as well as specifying the temporary installation directory for the .ini file.

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

Installing or upgrading automatically to a remote machine

1. Create a temporary installation directory on the primary site server. For example, create z:\temp_install.
2. Share the temporary installation directory.
3. Use the following steps if you downloaded your software from the web.
    a. Unzip the .exe file that you downloaded to another temporary directory.
    b. Locate the subdirectory under \setup\dt that is appropriate for your architecture, either i386 or x64.
    c. Copy the files from the \setup\dt\i386 or \setup\dt\x64 directory to your shared temporary installation directory.
4. Use the following steps if you have a DVD.
    a. Locate the subdirectory under \setup\dt that is appropriate for your architecture, either i386 or x64.
    b. Copy the files from the \setup\dt\i386 or \setup\dt\x64 directory to your shared temporary installation directory.
5. Remove the read-only attributes from the files in the shared temporary installation directory.
6. Make a backup copy of the default DTSetup.ini file in the shared temporary installation directory.
7. Edit DTSetup.ini as needed using the values described in the previous table.
8. From each server where you want to install Double-Take, map a drive to the shared temporary installation directory. For example, you might map your m: drive to the share.
9. Run the following case-sensitive command from the mapped drive.

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

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

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

```
C:\>net use m: \\server_name\share
The command completed successfully
C:\>M:
M:\>setup /s /v"DTSETUPINI=\"m:\DTSetup.ini\" /qn"
```

Installing using the Double-Take Console

You can use the Double-Take Console to install or upgrade Double-Take Availability on your other servers. The installation is a full, client and server installation.

If you are upgrading from Double-Take RecoverNow, you should manually remove it and any TimeData and OnTrack PowerControls components before upgrading. The uninstall may be time consuming while it removes the associated SQL instance and deletes the continuous data protection storage bins. Additionally, the TimeData uninstall may require a reboot. Make sure that you recall and restore any Double-Take archived or deduplicated files before you uninstall Double-Take RecoverNow. If you choose not to uninstall TimeData or Ontrack PowerControls, Double-Take RecoverNow will still be uninstalled, and these programs will no longer function after the upgrade. If you do not remove the products, you will have to manually remove them.

The push installation process will not work on cluster nodes.

If you are going to be creating a V to Hyper-V or V to ESX job (see *Selecting a protection type* on page 165 for details on which job to create), the job creation process can automatically push the Double-Take installation out to the servers, if you have the Double-Take Availability Virtual Guest Edition license. After the job is created, you will have 14 days to activate the licenses. (See *Activating a single license* on page 54 or *Activating multiple licenses* on page 56.) If you do not have a Double-Take Availability Virtual Guest Edition license, you will need to perform the installation outside of the job creation process.

---

1. Add the servers where you want to install Double-Take to your console session. See *Adding servers* on page 71.
2. From the **Manage Servers** page, highlight all of the servers where you want to install Double-Take, and select **Install** from the toolbar.
3. Each server needs an activation code for the installation. If you are upgrading and your server already has an activation code that is valid for the current version, you can either skip this step to use the existing activation code or you can complete this step to replace the existing activation code.
    a. If you have a single activation code that can be used on multiple servers, such as a site license or an evaluation code, highlight all of the servers where you want to apply the same activation code and click **Set Activation codes**. If you have unique activation codes for each of your servers, highlight just one server and click **Set Activation codes**.
    b. Type in your activation code or click **Choose from inventory** to open the Activation Codes dialog box where you can select the activation code you want to apply. For more information on the license inventory, see *Managing the Double-Take license inventory* on page 49.
    c. Click **OK** to configure the installation or upgrade for the selected activation code.
    d. Repeat steps a-c for any additional servers that are using unique activation codes.

If you are pushing out a code that requires activation, you will have 14 days to activate the code. See *License management and activation* on page 48 for more details on license activation.

4. The **Default Installation Options** section contains the default settings from the console's **Options** page. These settings will be applied to all of the servers you are installing to or upgrading. If desired, modify any of the installation options. See *Console options* on page 62.

Options you modify when using the push installation will be replaced by the default values from the **Options** page each time you use the push installation.

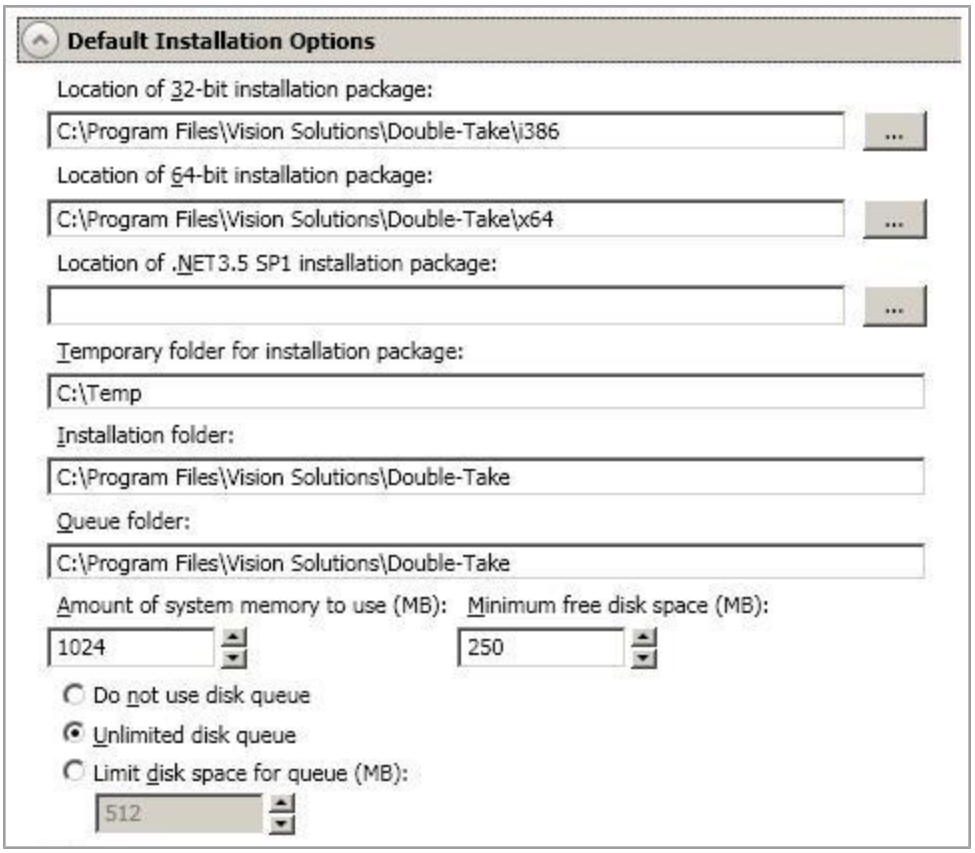

- **Location of 32-bit installation package**—Specify the location of the setup file (on the local machine) that will be used to install on 32-bit servers. By default, this is in the \i386 subdirectory where you installed Double-Take.
- **Location of 64-bit installation package**—Specify the location of the setup file (on the local machine) that will be used to install on 64-bit servers. By default, this is in the \x64 subdirectory where you installed Double-Take.
- **Location of .NET 3.5 SP1 installation package**—If your servers are running Windows 2008 or earlier and do not have Microsoft .NET version 3.5.1, specify the location of the setup file (on the local machine) that will be used to install it. The setup file is available on

the Double-Take DVD in the \NetFx\v3.5SP1\Full directory or from the Microsoft web site. For 2008 R2 or later, it will automatically be enabled for you through Windows features.

- **Temporary folder for installation package**—Specify a temporary location (on the server where you are installing Double-Take) where the installation files will be copied and run. You need approximately 130 MB of space in the specified location.
- **Installation folder**—Specify the location where you want to install Double-Take on each server.
- **Queue folder**—Specify the location where you want to store the Double-Take disk queue on each server.
- **Amount of system memory to use**—Specify the maximum amount of memory, in MB, that can be used for Double-Take processing. For complete details on memory usage, see *Double-Take queue* on page 88.
- **Minimum free disk space**—This is the minimum amount of disk space in the specified **Queue folder** that must be available at all times. This amount should be less than the amount of physical disk space minus the disk size specified for **Limit disk space for queue**.
- **Do not use disk queue**—This option will disable disk queuing. When system memory has been exhausted, Double-Take will automatically begin the auto-disconnect process.
- **Unlimited disk queue**—Double-Take will use an unlimited amount of disk space in the specified **Queue folder** for disk queuing, which will allow the queue usage to automatically expand whenever the available disk space expands. When the available disk space has been used, Double-Take will automatically begin the auto-disconnect process.
- **Limit disk space for queue**—This option will allow you to specify a fixed amount of disk space, in MB, in the specified **Queue folder** that can be used for Double-Take disk queuing. When the disk space limit is reached, Double-Take will automatically begin the auto-disconnect process.

5. Specify when you want to perform the installations under the **Schedule** section.

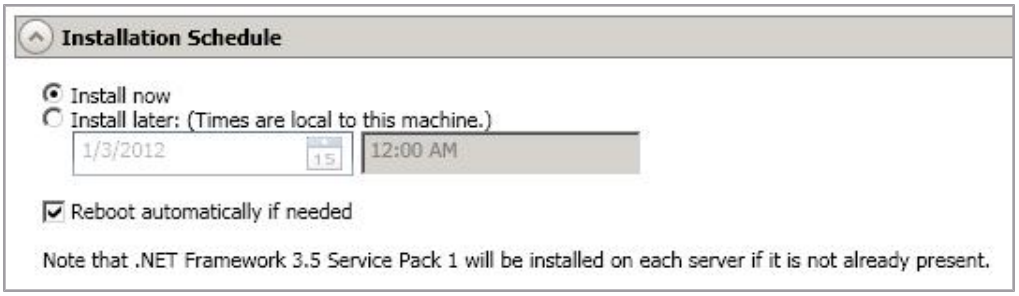

- **Install now**—Select this option to complete the installation immediately.
- **Install later**—Select this option and specify a date and time to complete the installation then.
- **Reboot automatically if needed**—If selected, the server will automatically reboot after the installation, if a reboot is required.

6. After you have configured your installation options, click **Install**.

During an upgrade, any existing jobs on the **Manage Jobs** page may disappear and then reappear. This is expected while certain Double-Take files are updated.

If there are errors with the push installation before the installation is executed on the server, check the console log on the machine where you are pushing the installation from. Once the installation execution has started, then check the installation log on the server where you are installing.

License management and activation

- **License management**—You can manage your Double-Take licenses through the license inventory feature in the Double-Take Console. Ideally, you should select one machine in your organization where you want to maintain the license inventory because the license inventory does not communicate between machines.

  From the license inventory, you can add and remove Double-Take licenses manually. You can import and export a Double-Take license file to handle groups of licenses. You can also activate and deactivate the licenses in your inventory. See *Managing the Double-Take license inventory* on page 49 for details on using the license inventory.

- **Licensing a server or appliance**—If you did not license your server during the installation (appliances are not licensed during the installation), you can apply a Double-Take license to a server or appliance using the console. See *Licensing a server* on page 51.
- **Activation**—For those Double-Take licenses that require activation, you have 14 days to activate the license. You can use the license inventory to activate multiple licenses at one time, or you can activate your licenses one at a time on a per server basis.
- **Deactivation**—Additionally, you can deactivate a Double-Take license if you uninstall Double-Take using the Double-Take Console. The license that was in use will be deactivated, and that license can be used on a new server. If you do not deactivate your license, you will have to complete a host transfer which will allow you to activate the license for use on another server.You can only complete two host transfers, and then the license becomes invalid. You can complete an unlimited number of deactivations. See *Deactivating licenses* on page 58.

Managing the Double-Take license inventory

You can manage your Double-Take licenses through the license inventory feature in the Double-Take Console. Ideally, you should select one machine in your organization where you want to maintain the license inventory because the license inventory does not communicate between machines.

By default, the license inventory feature is enabled. You can, and should, disable it on machines that are not maintaining your license inventory. To disable the license inventory, select **Options** from the console toolbar, deselect **Enable license inventory**, and click **Save**.

To manage your license inventory, select **Go**, **Manage License Inventory**.

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

You will see the following fields on the **Manage License Inventory** page.

Warning or error icon

Warnings indicate the license has not been activated, or it is temporary and will expire. Errors indicate the license has expired.

Serial Number

The serial number associated with the license

Product

The Double-Take product associated with the license

License Type

A short description of the type of license, including any expiration date or quantity information

Version

The product version number associated with the license

Server

The name of the server the license has been applied to, if any. If the same activation code is used on multiple servers, for example with evaluation licenses, you will see multiple entries in the license inventory for each server that is using that activation code.

Use the following toolbar controls to manage your licenses.

Import Licenses 

Imports all of the activation codes from a license inventory file into the license inventory. This is a file you may have received from Vision Solutions, or it may be from another Double-Take Console. Depending on the information contained in the file, you may be prompted to activate some of the servers in your license inventory.

Export Licenses 

Exports all of the activation codes in the license inventory to a license inventory file. This is a file that you may want to send to Vision Solutions when upgrading codes for a newer release or when activating or deactivating licenses. In this case go to https://activate.doubletake.com. You may also need to upload a file to https://activate.doubletake.com/hosttransferfile.aspx to complete host-transfers. You may also want this file so you can store and back up your activation codes.

Add Licenses 

Allows you to manually enter your activation codes. Enter the activation codes in the Add Licenses dialog box in the space provided, separating multiple codes by a comma or by putting each code on a separate line. Once your activation codes are in your license inventory, you can apply them to a server in several ways. See *Activating a single license* on page 54 or See *Activating multiple licenses* on page 56.

Remove License 

Removes the selected activation code from the license inventory. You can only remove activation codes that are not being used by any server in the console.

Reclaim License 

Returns a **Single** license type back into the console's license inventory. You may want to reclaim a single license if you have removed a license from a server so that you can assign it to another server. This process should not be used for transferring an **Activated** license to a new or rebuilt server, which is considered deactivation. See *Deactivating licenses* on page 58 for details on deactivating a license.

Activate 

Activates and deactivates the activation codes in the license inventory. For complete details on this process, see *Activating multiple licenses* on page 56 and *Deactivating licenses* on page 58.

Licensing a server

If you did not license your server during the installation, you can apply a Double-Take license to a server using the console.

1. Make sure you have your server inserted in the console.
2. From the **Manage Servers** page, double-click on the server to view the server's details.
3. From the **View Server Details** page, click on the **Edit server properties** link.
4. Expand the **Licensing** section.
5. Licensing identifies your Double-Take activation codes.

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

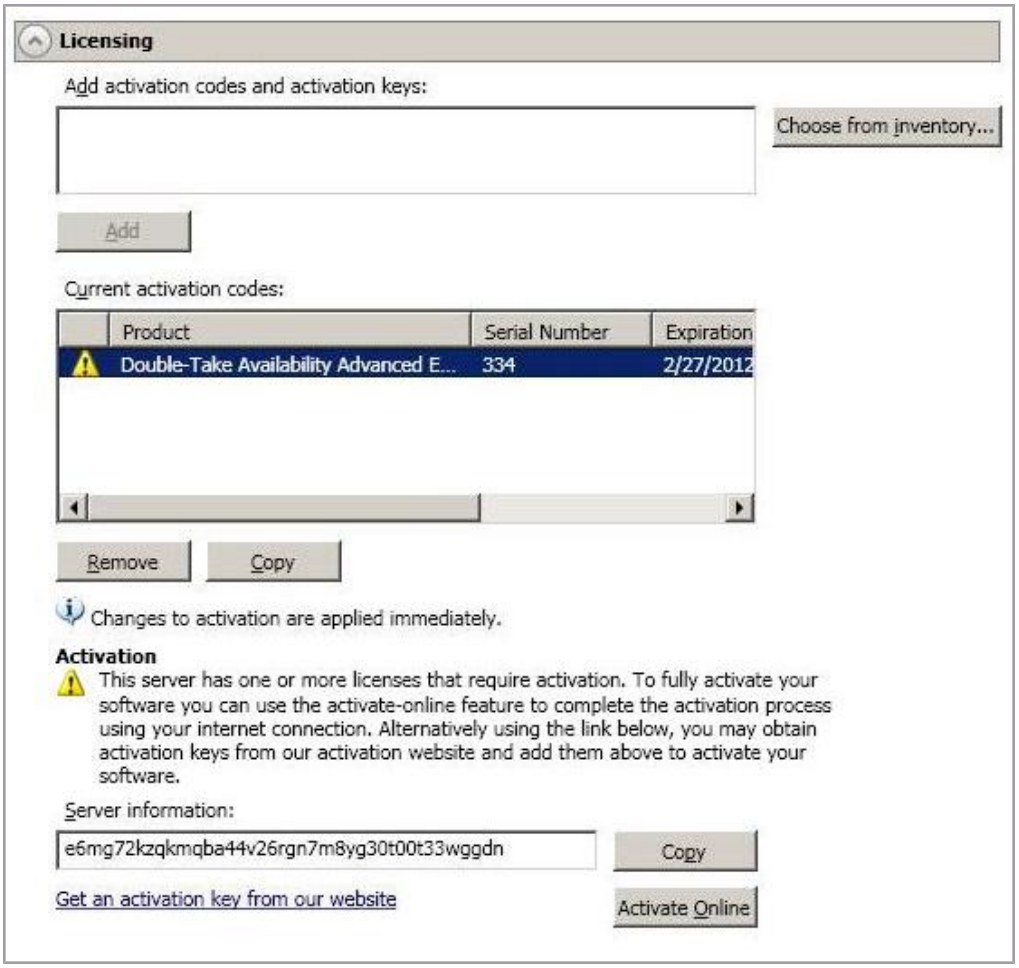

- **Add activation codes and activation keys**—The activation code and activation key are the Double-Take license which is required on every Double-Take server. They are a 24 character, alpha-numeric code. You can change your activation code without reinstalling, if

your license changes.

There are different licenses available.

- **Evaluation**—An evaluation license has an expiration date built into the activation code. When the license expires, the software will no longer function. The same evaluation license can be used on multiple machines on a network. An evaluation license does not have to be activated.
- **Single**—A single license is available on a per-machine basis. Each server is required to have a unique license whether it is functioning as a source, target, or both. A single license can only be used on one server. It may have a 14 day activation period built-in.
- **Site**—A site license is available to register every machine with the same license. This license is designed to be used on multiple servers on a network. This license may or may not have to be activated, depending on the code.

To add an activation code and activation key, type in the code and click **Add**. If your console has been enabled to manage your license inventory, click **Choose from inventory** to open the Activation Codes dialog box where you can select the activation codes you want to apply. See *Console options* on page 62 for details on enabling the license inventory.

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

- **Current activation codes**—The server's current activation codes are displayed.
  - **Warning or error icon**—Warnings indicate the license is temporary and will expire. Errors indicate the license has expired.
  - **Product**—The product associated with the license
  - **Serial Number**—The serial number associated with the license
  - **Expiration Date**—The date the license expires, if there is one
  - **Activation Code**—The activation code

  To remove a code, highlight it and click **Remove**. To copy a code, highlight it and click **Copy**.
- **Activation**—If your activation code needs to be activated, you will see an additional **Activation** section at the bottom of the **Licensing** section. To activate your code, use one of the following procedures.
  - **Activate online**—If you have Internet access, you can activate your license and apply the activated license to the server in one step. Select **Activate Online**. You will not be able to activate a license that has already been activated but has not been deactivated. In that case, you will be prompted to complete a host transfer. Ideally, you should deactivate the license instead of doing a host transfer. See *Deactivating licenses* on page 58 for more details. If you must do a host transfer, you will have to export your license inventory to a file, upload that file to https://activate.doubletake.com/hosttransferfile.aspx, and then import the file you get back. See *Managing the Double-Take license inventory* on page 49 for details on

exporting and importing a license inventory file.

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

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

If your activation codes needs to be activated, you will have 14 days to do so.

If you 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). See *Deactivating licenses* on page 58. 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.

- **Total licenses quantity**—If your server is a Double-Take controller appliance, you will see the available and total license count displayed. The controller appliance handles all license management for your agentless vSphere jobs. Each agentless vSphere job will use one license from your available quantity. If you create jobs for five virtual machines at one time, five licenses will be used. When you delete a job, the license associated with the job will be released to be available for another job. Stopping a job will not delete the license associated with the job. If you have no additional licenses available, you cannot create any new jobs. You can add to your available license quantity by purchasing additional licenses.

6. Once you have completed your licensing, click **OK** to return to the **Manage Servers** page.

Activating a single license

You will need to have Double-Take installed on a single server, then you can activate the license on that server using the Double-Take Console. You can repeat this process for other servers individually.

1. Click **Manage Servers** in the toolbar.
2. Click **Add Servers** in the **Manage Servers** page toolbar.
3. Specify the name of your server and the credentials of the account used to install Double-Take, and click **Add**.
4. Click **OK**.
5. After you server has been added to the **Manage Servers** page, click **View Server Details** in the toolbar.
6. On the **View Server Details** page, click **Edit server properties** under **Tasks**.
7. Expand the **Licensing** section.

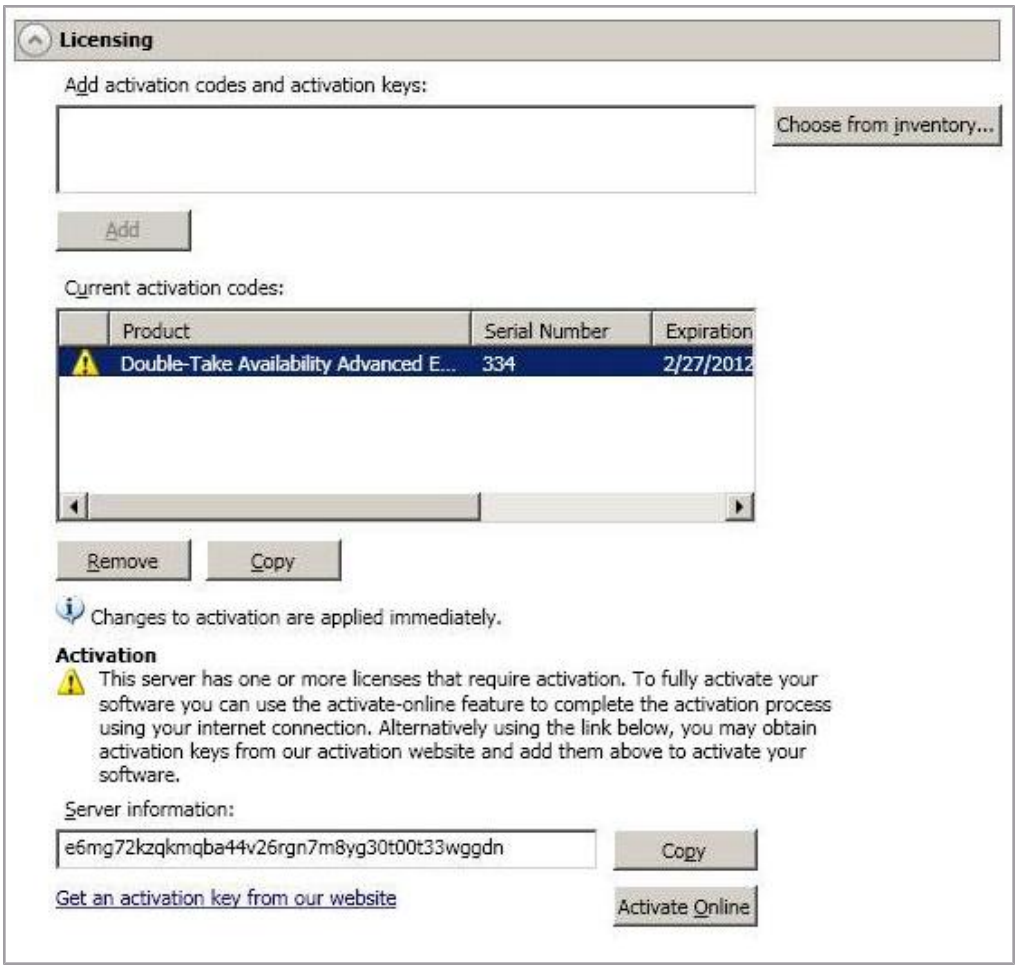

8. If your activation code needs to be activated, you will see an additional **Activation** section at the bottom of the **Licensing** section. To activate your code, use one of the following procedures.
   - **Activate online**—If you have Internet access, you can activate your license and apply the activated license to the server in one step. Select **Activate Online**. You will not be able to

activate a license that has already been activated but has not been deactivated. In that case, you will be prompted to complete a host transfer. Ideally, you should deactivate the license instead of doing a host transfer. See *Deactivating licenses* on page 58 for more details. If you must do a host transfer, you will have to export your license inventory to a file, upload that file to https://activate.doubletake.com/hosttransferfile.aspx, and then import the file you get back. See *Managing the Double-Take license inventory* on page 49 for details on exporting and importing a license inventory file.

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

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

If your activation codes needs to be activated, you will have 14 days to do so.

If you 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). See *Deactivating licenses* on page 58. 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.

Activating multiple licenses

You will need to have Double-Take installed on multiple servers, then you can activate the licenses on those servers all at once.

1. On the **Manage Servers** page, verify that all of the Double-Take servers that you want to activate have been added to your Double-Take Console. See *Adding servers* on page 71.
2. Select **Go**, **Manage License Inventory**.
3. Click **Activate** in the toolbar.

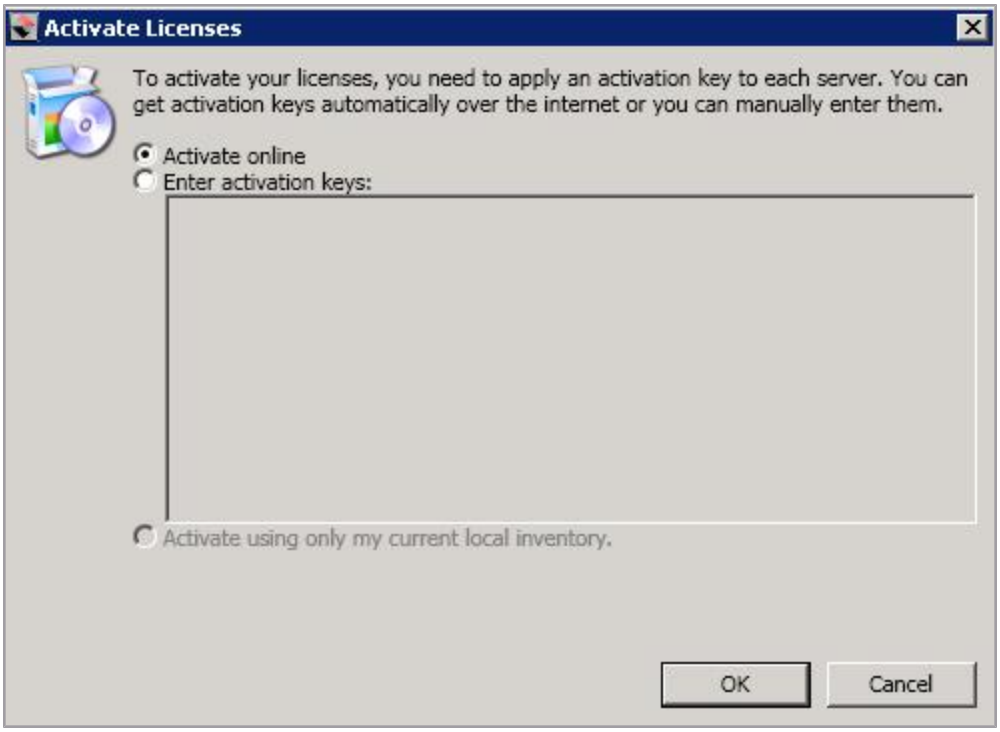

- **Activate online**—Select this option on an Internet connected machine to contact the Vision Solutions activation web site to automatically activate all of the licenses in your license inventory. This process will also deactivate any licenses that have been uninstalled through the Double-Take Console. You will not be able to activate a license that has already been activated but has not been deactivated. In that case, you will be prompted to complete a host transfer. Ideally, you should deactivate the license instead of doing a host transfer. See *Deactivating licenses* on page 58 for more details. If you must do a host transfer, you will have to export your license inventory to a file, upload that file to https://activate.doubletake.com/hosttransferfile.aspx, and then import the file you get back. See *Managing the Double-Take license inventory* on page 49 for details on exporting and importing a license inventory file.
- **Enter activation keys**—Select this option on a non-Internet connected machine to manually enter activation keys. You can obtain these keys from https://activate.doubletake.com by entering the server information for each server manually or by uploading an export file of your license inventory. See *Activating a single license* on page 54 for details on gathering the server information or *Managing the Double-Take license inventory* on page 49 for details on creating an export file.

- **Activate using only my current local inventory**—Select this option to activate all of the licenses in your inventory based on the inventory's current local settings. This process will also deactivate any licenses that have been uninstalled through the Double-Take Console. For example, if a server was offline when you performed activate online, you can perform the activation process again when the server is available. You may also need to perform this option if you selected not to activate the servers when you imported a license inventory file.

4. Click **OK** to begin the activation process.

After your servers have been activated, you will have permanent codes that are specific to each server. They cannot be used on any other server. If the activated code and server do not match, Double-Take will not run.

If your activation codes needs to be activated, you will have 14 days to do so.

If you 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). See *Deactivating licenses* on page 58. 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.

Deactivating licenses

Once you have activated a license, it is a permanent code specific to the assigned server. It cannot be used on any other server. If the activation code and server do not match, Double-Take will not run. You can deactivate a Double-Take license if you uninstall Double-Take using the Double-Take Console. The license that was in use will be deactivated, and that license can be used on a new server. If you do not deactivate your license, you will have to complete a host transfer which will allow you to activate the license for use on another server.You can only complete two host transfers, and then the license becomes invalid. You can complete an unlimited number of deactivations.

The key to the deactivation process is performing the uninstall through the Double-Take Console. If you perform the uninstall through the Windows add remove or programs and features applet, the license will not be deactivated.

If you failed over a full server job but did not enable reverse protection when you configured the job, you will lose the activated target license. To workaround this situation, uninstall Double-Take on the failed over source (currently on the target hardware) using the Double-Take Console, and the uninstall process will deactivate both the source and target licenses. You can then reinstall Double-Take on the failed over source (currently on the target hardware) and reactivate the source's original license.

---

1. From the **Manage Jobs** page, confirm that there are no jobs currently using the server that you want to uninstall from. If there are, stop and delete those jobs.
2. From the **Manage Servers** page, highlight the server you want to uninstall from and click **Uninstall** from the toolbar.
3. When the uninstallation is complete, select **Go**, **Manage License Inventory**.
4. Click **Activate** in the toolbar.
5. Select **Activate online**. This option will contact the Vision Solutions activation web site to automatically deactivate the licenses that have been uninstalled through the Double-Take Console.
6. Click **OK** to begin the process.

After your licenses have been deactivated, you can reuse them on another Double-Take server.

If your console is not connected to the Internet, you will have to export your license inventory to a file and then upload that file to https://activate.doubletake.com to complete the deactivation process. See *Managing the Double-Take license inventory* on page 49 for details on export the license inventory.

Chapter 4 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 **Home** page. This page provides a high-level overview of the status of your jobs.

The appearance of the **Home** page is the same for all users. However, other console pages may have variances in the appearance depending on the Double-Take products that you have installed, the Double-Take activation codes on your servers, and the type of job you are working with.

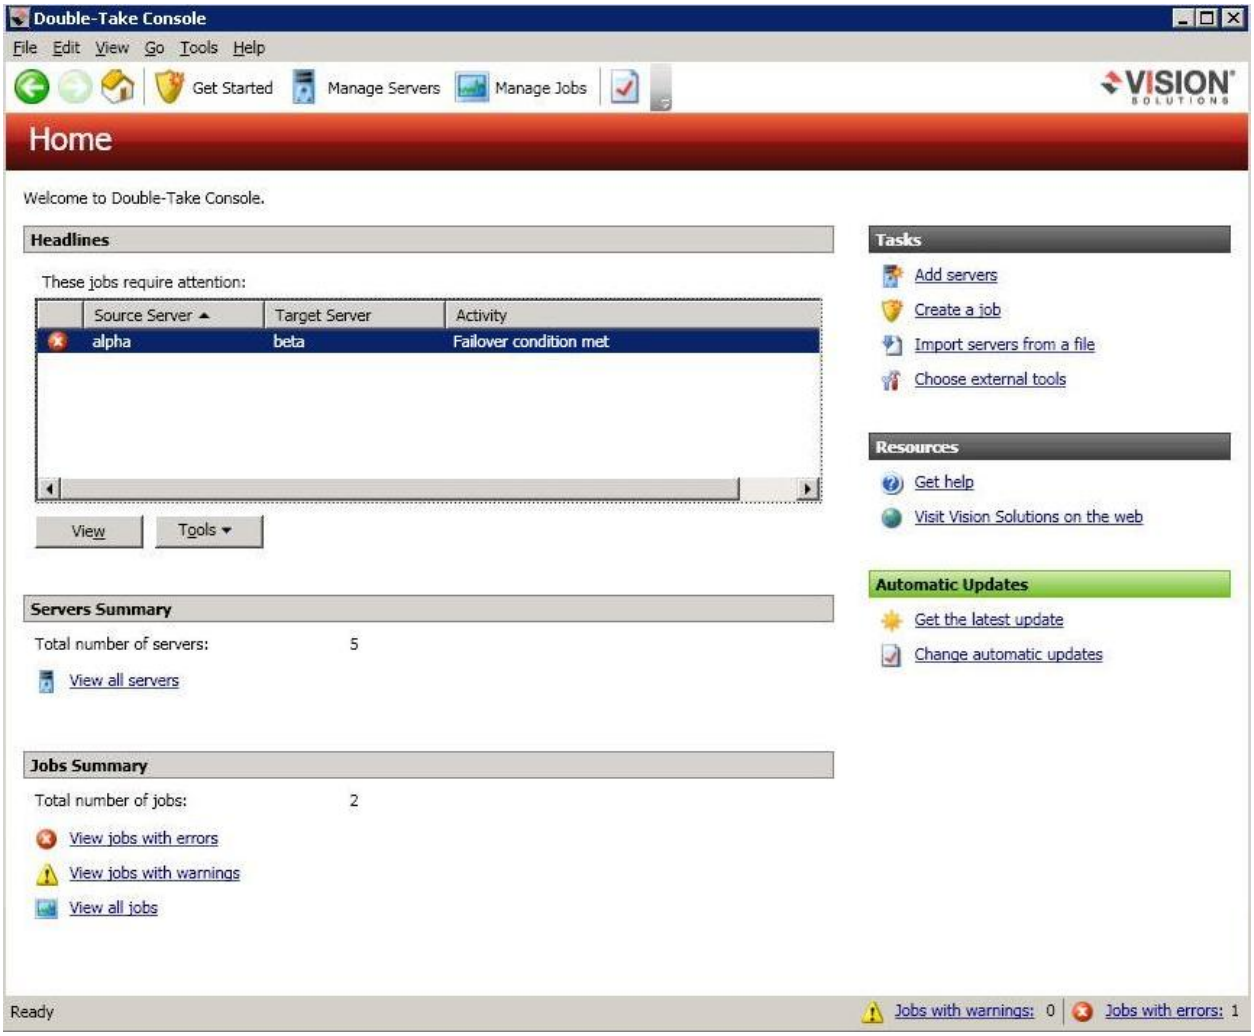

- **Headlines**—The top section gives a quick overview of any jobs that require attention as well as providing quick access buttons.
  - **These jobs require attention**—Any jobs that require attention (those in an error state) are listed. You will see the source and target server names listed, as well as a short description of the issue that requires your attention. If the list is blank, there are no jobs that

require immediate attention.

  - **View**—If you highlight a job in the list and click **View**, you will go to the **View Job Details** page where you can see more detailed information about the job.
  - **Tools**—Select this drop-down list to launch other Vision Solutions consoles.
- **Servers Summary**—The middle section summarizes the servers in your console.
  - **Total number of servers**—This field displays the number of servers that you have been added to the console.
  - **View all servers**—Select this link to go to the **Manage Servers** page where you can view, edit, add, remove, or manage the servers in your console. See *Managing servers* on page 65.
- **Jobs Summary**—The bottom section summarizes the jobs in your console.
  - **Total number of jobs**—This field displays the number of jobs running on the servers in your console.
  - **View jobs with errors**—Select this link to go to the **Manage Jobs** page, where the **Filter: Jobs with errors** will automatically be applied.
  - **View jobs with warnings**—Select this link to go to the **Manage Jobs** page, where the **Filter: Jobs with warnings** will automatically be applied.
  - **View all jobs**—Select this link to go to the **Manage Jobs** page and view all jobs.

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 **Manage Jobs** page, where the appropriate **Filter: Jobs with warnings** or **Filter: Jobs with errors** will automatically be applied.

Double-Take Console requirements

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

- **Operating system**—The Double-Take Console can be run from a Windows source or target. It can also be run from a 32-bit or 64-bit physical or virtual machine running Windows 8, Windows 7, Windows Vista, or Windows XP Service Pack 2 or later.
- **Microsoft .NET Framework**—Microsoft .NET Framework version 3.5 Service Pack 1 is required. This version is not included in the .NET version 4.0 release. Therefore, even if you have .NET version 4.0 installed, you will also need version 3.5.1. For Windows 2008 and earlier, you can install this version from the Double-Take DVD, via a web connection during the Double-Take installation, or from a copy you have obtained manually from the Microsoft web site. For Windows 2008 R2 and later, you need to enable it through Windows features.
- **Screen resolution**—For best results, use a 1024x768 or higher screen resolution.

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

Console options

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

- **Monitoring interval**—Specifies how often, in seconds, the console refreshes the monitoring data. The servers will be polled at the specified interval for information to refresh the console.
- **Automatic retry**—This option will have the console automatically retry server login credentials, after the specified retry interval, if the server login credentials are not accepted. Keep in mind the following caveats when using this option.
  - This is only for server credentials, not job credentials.
  - A set of credentials provided for or used by multiple servers will not be retried for the specified retry interval on any server if it fails on any of the servers using it.
  - Verify your environment's security policy when using this option. Check your policies for failed login lock outs and resets. For example, if your policy is to reset the failed login attempt count after 30 minutes, set this auto-retry option to the same or a slightly larger value as the 30 minute security policy to decrease the chance of a lockout.
  - Restarting the Double-Take Console will automatically initiate an immediate login.
  - Entering new credentials will initiate an immediate login using the new credentials.
- **Retry on this interval**—If you have enabled the automatic retry, specify the length of time, in minutes, to retry the login.
- **Default port for XML web services protocol**—Specifies the port that the console will use when sending and receiving data to Double-Take servers. By default, the port is 6325. Changes to the console port will not take effect until the console is restarted.
- **Default port for legacy protocol**—If you are using an older Double-Take version, you will need to use the legacy protocol port. This applies to Double-Take versions 5.1 or earlier.
- **Export Diagnostic Data**—This button creates a raw data file that can be used for debugging errors in the Double-Take Console. Use this button as directed by technical support.
- **View Log File**—This button opens the Double-Take Console log file. Use this button as directed by technical support.
- **View Data File**—This button opens the Double-Take Console data file. Use this button as directed by technical support.
- **Automatically check for updates**—By default, each time the console is started, it will automatically check the Vision Solutions web site to see if there is updated console software available. If there is updated console software available, an **Automatic Updates** section will appear on the **Home** page. Click **Get the latest update** to download and install the updated console software.

  If you want to disable the automatic check for updates, click **Change automatic updates** or select **Options** from the toolbar. On the **Options** page, deselect **Automatically check for updates** to disable the automatic check.

  You can also manually check for updates by selecting **Help**, **Check for Updates**.

  - **Update available**—If there is an update available, click **Get Update**. The dialog box will close and your web browser will open to the Vision Solutions web site where you can download and install the update.

- **No update available**—If you are using the most recent console software, that will be indicated. Click **Close**.
- **No connection available**—If the console cannot contact the update server of if there is an error, the console will report that information. The console log contains a more detailed explanation of the error. Click **Check using Browser** if you want to open your browser to check for console software updates. You will need to use your browser if your Internet access is through a proxy server.

- **Enable license inventory**—This option allows you to use this console to manage the Double-Take licenses assigned to your organization. When this option is enabled, the **Manage License Inventory** page is also enabled. See *Managing the Double-Take license inventory* on page 49.

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

- **Default Installation Options**—All of the fields under the **Default Installation Options** section are used by the push installation on the **Install** page. See *Installing using the Double-Take Console* on page 44. The values specified here will be the default options used for the push installation. Options you modify when using the push installation will be replaced by these default values each time you use the push installation.
  - **Location of 32-bit installation package**—Specify the location of the setup file (on the local machine) that will be used to install on 32-bit servers. By default, this is in the \i386 subdirectory where you installed Double-Take.
  - **Location of 64-bit installation package**—Specify the location of the setup file (on the local machine) that will be used to install on 64-bit servers. By default, this is in the \x64 subdirectory where you installed Double-Take.
  - **Location of .NET 3.5 SP1 installation package**—If your servers are running Windows 2008 or earlier and do not have Microsoft .NET version 3.5.1, specify the location of the setup file (on the local machine) that will be used to install it. The setup file is available on the Double-Take DVD in the \NetFx\v3.5SP1\Full directory or from the Microsoft web site. For 2008 R2 or later, it will automatically be enabled for you through Windows features.
  - **Temporary folder for installation package**—Specify a temporary location (on the server where you are installing Double-Take) where the installation files will be copied and run. You need approximately 130 MB of space in the specified location.
  - **Installation folder**—Specify the location where you want to install Double-Take on each server.
  - **Queue folder**—Specify the location where you want to store the Double-Take disk queue on each server.
  - **Amount of system memory to use**—Specify the maximum amount of memory, in MB, that can be used for Double-Take processing. For complete details on memory usage, see *Double-Take queue* on page 88.
  - **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.

Chapter 5 Managing servers

To manage the servers in your console, select **Manage Servers** from the toolbar. The **Manage Servers** page allows you to view, edit, add, remove, or manage the servers in your console.

You can also organize the servers that are in your console into groups, allowing you to filter the servers you are viewing based on your organization. The servers displayed in the right pane depend on the server group folder selected in the left pane. Every server in your console session is displayed when the **All Servers** group is selected. If you have created and populated server groups under **My Servers**, then only the servers in the selected group will displayed in the right pane.

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

Right pane display

The following table identifies the columns displayed in the right pane of the **Manage Servers** page.

Column 1 (Blank)

The first blank column indicates the machine type.

 Double-Take source or target server which could be a physical server, virtual machine, or a cluster node

 Double-Take source or target server which is a Windows cluster

 VMware server which could be a vCenter server or an ESX or ESXi host.

 Double-Take controller appliance

 Double-Take replication appliance

 Double-Take Reporting Service server

 Offline server which means the console cannot communicate with this machine.

 Server error which means the console can communicate with the machine, but it cannot communicate with Double-Take on it.

Column 2 (Blank)

The second blank column indicates the security level

Processing—The console is attempting to communicate with machine.

Administrator access—This level grants full control.

Monitor only access—This level grants monitoring privileges only.

No security access—This level does not allow monitoring or control.

Server

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

Activity

There are many different **Activity** messages that keep you informed of the server activity. Most of the activity messages are informational and do not require any administrator interaction. If you see error messages, check the server details. See *Viewing server details* on page 78.

Version

The product version information

Licensing Status

The status of the license on the server. If your license is expired, any jobs using that server will be in an error state.

Product

The Double-Take products licensed for the server or the Double-Take role for the server.

Activation Code

The activation codes associated with the products licensed for the server. If your license is not valid for the operating system on your server, the activation code will be identified as Invalid Activation Code. There will be no activation code listed for those servers that are not licensed, like a VMware server.

Serial Number

The serial number associated with the activation code

Main toolbar and right-click menu

The following options are available on the main toolbar of the **Manage Servers** page and the right-click menu. Some of the options are only available in the right-click menu. Each of the options control the server that is selected in the right pane.

Add Servers 

Adds a new server. This button leaves the **Manage Servers** page and opens the **Add Servers** page. See *Adding servers* on page 71.

Add Replication Appliance 

Adds a new replication appliance. This option is only valid for agentless vSphere protection.

View Server Details 

Views detailed information about a server. This button leaves the **Manage Servers** page and opens the **View Server Details** page. See *Viewing server details* on page 78.

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 77. You will remain on the **Manage 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 a full server job with reverse protection enabled, you need to update the target image stored on the source if you change the credentials on the target server. See *Viewing full server job details* on page 273.

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. See *Installing using the Double-Take Console* on page 44.

Uninstall 

Uninstalls Double-Take on the selected server.

View Server Events 

Views event messages for a server. This button leaves the **Manage Servers** page and opens the **View Server Events** page. See *Viewing server events* on page 157.

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. See *Viewing the log files through the Double-Take Console* on page 678 for more details on this view.

Refresh

Refreshes the status of the selected servers.

Gather Support Diagnostics

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

View Replication Service Details

Views the replication service details for a server. This 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 *Replication service view* on page 697 for more details on this view.

Overflow Chevron 

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

Left pane toolbar

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.

Adding servers

The first time you start the console, the **Manage Servers** page is empty. In order to protect and monitor your servers, you must insert your servers and/or appliances in the console. For some jobs, you can insert servers and appliances during job creation, or you have three other methods for inserting servers into the console.

If you will be creating a full server to ESX appliance job (see *Selecting a protection type* on page 165 for details on selecting a job type), you should insert your virtual recovery appliance into your console session before inserting your source servers because a source server needs to be associated with an appliance that is already inserted in the console.

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. See the following NAT configuration section if you have a NAT environment.
   - **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.
   - **Password**—Specify the password associated with the **User name** you entered.
4. You can expand the **More Options** section to configure the following settings.
   - **Domain**—If you are working in a domain environment, specify the **Domain**.
   - **Associate with Double-Take Linux Appliance**—If you will be creating a full server to ESX appliance job, specify the name or IP address of your virtual recovery appliance. Your Linux source is now considered a proxied server.
5. After you have specified the server or appliance information, click **Add**.
6. Repeat steps 3 through 5 for any other servers or appliances you want to add.
7. 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.
8. When your list of **Servers to be added** is complete, click **OK**.

NAT configuration

If you are going to create a files and folders (non-cluster) job, a full server job without reverse protection, or a full server to ESX appliance job, then your servers can be in a NAT environment. Other job types do not support a NAT environment.

The name or IP address you use to add a server to the console is dependent on where are you are running the console. It is also dependent on the job type you are using. If you are using a full server to ESX appliance job, the console only communicates with the appliance. Therefore, you need to use the address which the appliance has access to when adding the source server. Use the one of the following tables to determine what name or IP address to use when adding a server or appliance to the console.

In these tables, public addresses are those addresses that are publicly available when a server is behind a NAT router. Private addresses are those addresses that are privately available when a server is behind a NAT router. An address that is not labeled as public or private are for servers that are not behind a NAT router. This is generally a public address but is not named as such in these tables to try to more clearly identify when a public NAT address needs to be used.

Files and folders (non-cluster) jobs and full server jobs without reverse protection

Location of servers	Location of Double-Take Console	How to add the server to the Double-Take Console
If your source and target are behind individual NAT routers,	and your Double-Take Console is located behind the NAT router with the source,	specify the name or private IP address of the source and the public IP address of the target (which is the public IP address of the target's NAT router).
	and your Double-Take Console is located behind the NAT router with the target,	specify the public IP address of the source (which is the public IP address of the source's NAT router) and the name or private IP address of the target.
	and your Double-Take Console is located between the two NAT routers,	specify the public IP address of the source (which is the public IP address of the source's NAT router) and the public IP address of the target (which is the public IP address of the target's NAT router).
	and your Double-Take Console is located behind a third NAT router,	

If your source is behind a NAT router but your target is not,	and your Double-Take Console is located behind the NAT router with the source,	specify the name or private IP address of the source and the name or IP address of the target.
	and your Double-Take Console is located on the target network,	specify the public IP address of the source (which is the public IP address of the source's NAT router) and the name or IP address of the target.
If your target is behind a NAT router but your source is not,	and your Double-Take Console is located behind the NAT router with the target,	specify the name or IP address of the source and the name or private IP address of the target.
	and your Double-Take Console is located on the source network,	specify the name or IP address of the source and the public IP address of the target (which is the public address of the target's NAT router).
If your source and target are both behind a single NAT router with multiple public NICs,	and your Double-Take Console is located outside of the router,	specify the public IP addresses for the source and the target.

Full server to ESX appliance jobs

Location of source and appliance	Location of Double-Take Console	How to add the source and appliance to the Double-Take Console
If your source and appliance are behind individual NAT routers,	and your Double-Take Console is located behind the NAT router with the source,	specify the public IP address of the source (which is the public IP address of the source's NAT router) and the public IP address of the appliance (which is the public IP address of the appliance's NAT router).
	and your Double-Take Console is located behind the NAT router with the appliance,	specify the public IP address of the source (which is the public IP address of the source's NAT router) and the name or private IP address of the appliance.
	and your Double-Take Console is located between the two NAT routers,	specify the public IP address of the source (which is the public IP address of the source's NAT router) and the public IP address of the appliance (which is the public IP address of the appliance's NAT router).
	and your Double-Take Console is located behind a third NAT router,	
If your source is behind a NAT router but your appliance is not,	and your Double-Take Console is located behind the NAT router with the source,	specify the public IP address of the source (which is the public IP address of the source's NAT router) and the name or IP address of the appliance.
	and your Double-Take Console is located on the appliance network,	specify the public IP address of the source (which is the public IP address of the source's NAT router) and the name or IP address of the appliance.
If your appliance is behind a NAT router but your source is not,	and your Double-Take Console is located behind the NAT router with the appliance,	specify the name or IP address of the source and the name or private IP address of the appliance.
	and your Double-Take Console is located on the source network,	specify the name or IP address of the source and the public IP address of the appliance (which is the public address of the appliance's NAT router).
If your source and appliance are both behind a single NAT router,	and your Double-Take Console is located outside of the router,	specify the name or private IP address of the source and the public IP address of the appliance (which is the public address of the apliance's NAT router).

As noted above, make sure you insert your virtual recovery appliance into your console session before inserting your source servers because a source server needs to be associated with an appliance that is already inserted in the console.

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 **Manage 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 **Manage Servers** page. When prompted, specify the **User name**, **Password**, and **Domain** of the account you want to use for this server. Click **OK** to save the changes.

Viewing server details

Highlight a server on the **Manage 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.

- **EngineRole**—Source or target server
- **ProxyRole**—A Linux appliance for a full server to ESX appliance job
- **ProxiedRole**—A Linux source server for a full server to ESX appliance job
- **ControllerRole**—Controller appliance for an agentless vSphere job
- **ReplicationApplianceRole**—Replication appliance for an agentless vSphere job
- **Reporting Service**—Double-Take Reporting Service server

Status

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

Activity

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

Connected via

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

Version

The product version information

Access

The security level granted to the specified user

User name

The user account used to access the server

Licensing

Licensing information for the server

- **Warning or error icon**—Warnings indicate the license is temporary and will expire. Errors indicate the license has expired or it is invalid for the operating system on the server.
- **Product**—The product associated with the license
- **Serial Number**—The serial number associated with the license
- **Expiration Date**—The date the license expires, if there is one
- **Activation Code**—The activation code
- **Licensing Status**—The status of the license on the server. If your license is expired, any jobs using that server will be in an error state.
- **License Type**—The type of Double-Take license

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.

Editing server properties

Highlight a server on the **Manage Servers** page and click **View Server Details** from the toolbar. Under **Tasks**, select **Edit server properties**. The **Edit Server Properties** page allows you to view and edit properties for that server. Click on a heading on the **Edit Server Properties** page to expand or collapse a section of properties.

- *General server properties* on page 81—Identifies the server
- *Server licensing* on page 82—Views, adds, and removes activation codes
- *Server setup properties* on page 85—Indicates how the server will act on startup and shutdown
- *Double-Take queue* on page 88—Configures the Double-Take queues
- *Source server properties* on page 92—Configures the source server
- *Target server properties* on page 95—Configures the target server
- *E-mail notification configuration* on page 97—Configures e-mail notification
- *Script credentials* on page 99—Specifies credentials to be used when executing custom scripts during mirroring or failover
- *Log file properties* on page 100—Configures log files

Server properties cannot be edited on a cluster.

General server properties

The general server properties identify the server.

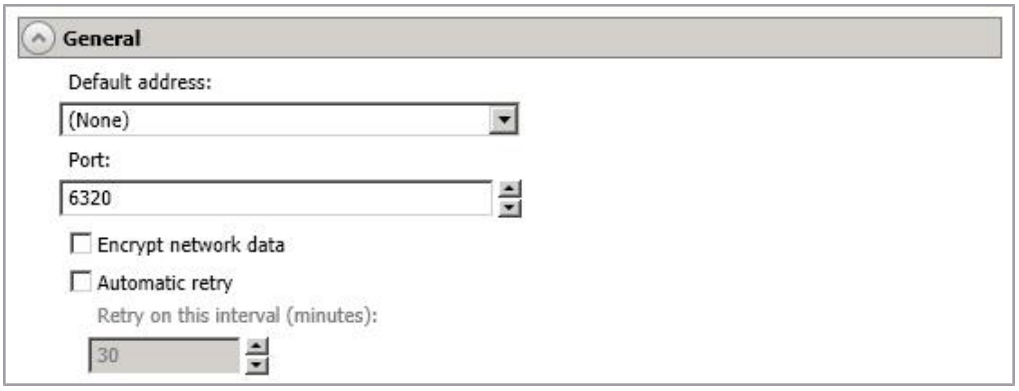

- **Default address**—On a server with multiple NICs, you can specify which address Double-Take traffic will use. It can also be used on servers with multiple IP addresses on a single NIC. If you change this setting, you must restart the Double-Take service for this change to take effect.
- **Port**—The server uses this port to send and receive commands and operations between Double-Take servers.
- **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 activation codes.

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

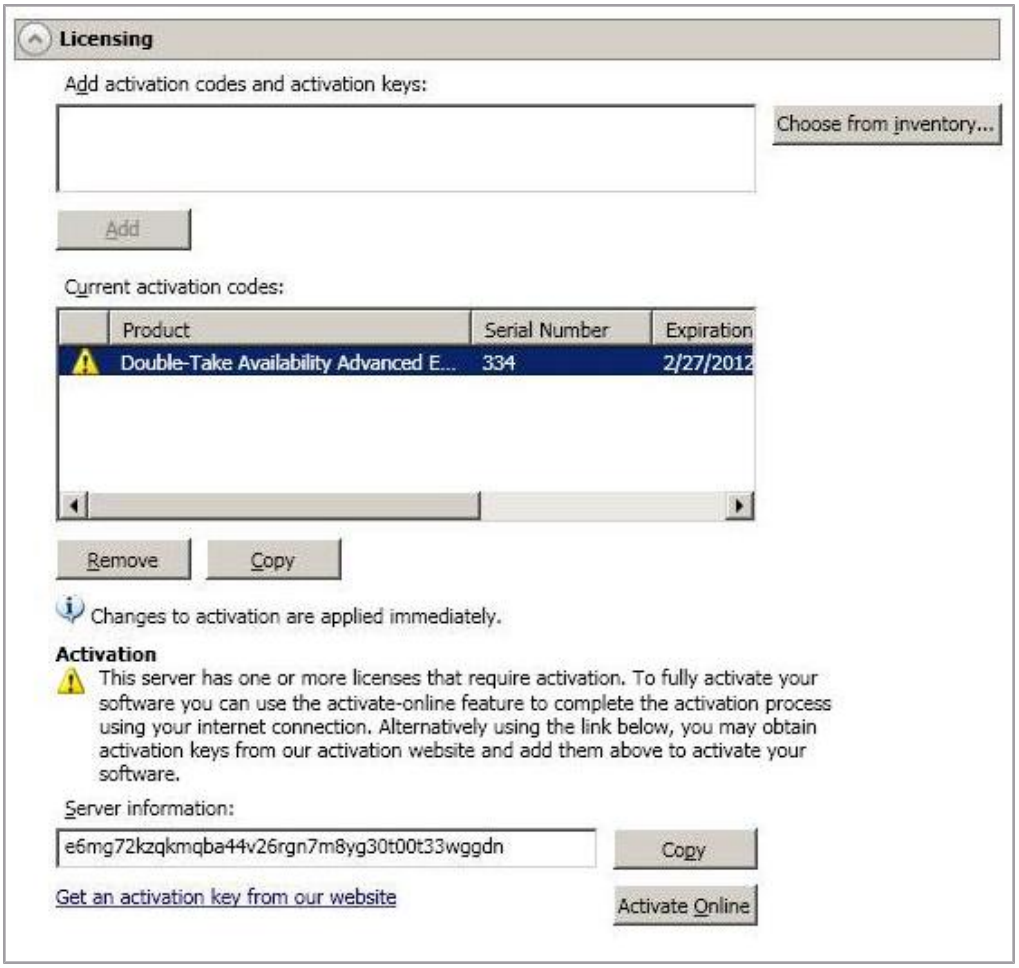

- **Add activation codes and activation keys**—The activation code and activation key are the Double-Take license which is required on every Double-Take server. They are a 24 character, alpha-numeric code. You can change your activation code without reinstalling, if your license changes.

  There are different licenses available.

  - **Evaluation**—An evaluation license has an expiration date built into the activation code. When the license expires, the software will no longer function. The same evaluation license can be used on multiple machines on a network. An evaluation license does not have to be activated.

- **Single**—A single license is available on a per-machine basis. Each server is required to have a unique license whether it is functioning as a source, target, or both. A single license can only be used on one server. It may have a 14 day activation period built-in.
- **Site**—A site license is available to register every machine with the same license. This license is designed to be used on multiple servers on a network. This license may or may not have to be activated, depending on the code.

To add an activation code and activation key, type in the code and click **Add**. If your console has been enabled to manage your license inventory, click **Choose from inventory** to open the Activation Codes dialog box where you can select the activation codes you want to apply. See *Console options* on page 62 for details on enabling the license inventory.

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

- **Current activation codes**—The server's current activation codes are displayed.
  - **Warning or error icon**—Warnings indicate the license is temporary and will expire. Errors indicate the license has expired.
  - **Product**—The product associated with the license
  - **Serial Number**—The serial number associated with the license
  - **Expiration Date**—The date the license expires, if there is one
  - **Activation Code**—The activation code

  To remove a code, highlight it and click **Remove**. To copy a code, highlight it and click **Copy**.
- **Activation**—If your activation code needs to be activated, you will see an additional **Activation** section at the bottom of the **Licensing** section. To activate your code, use one of the following procedures.
  - **Activate online**—If you have Internet access, you can activate your license and apply the activated license to the server in one step. Select **Activate Online**. You will not be able to activate a license that has already been activated but has not been deactivated. In that case, you will be prompted to complete a host transfer. Ideally, you should deactivate the license instead of doing a host transfer. See *Deactivating licenses* on page 58 for more details. If you must do a host transfer, you will have to export your license inventory to a file, upload that file to https://activate.doubletake.com/hosttransferfile.aspx, and then import the file you get back. See *Managing the Double-Take license inventory* on page 49 for details on exporting and importing a license inventory file.
  - **Obtain activation key online, then activate**—If you have Internet access, click the hyperlink in the **Activation** section to take you to the web so that you can submit your activation information. Complete and submit the activation form, and you will receive an e-mail with the activation key. Activate your server by entering the activation key in the **Add activation codes and activations keys** field and clicking **Add**.
  - **Obtain activation key offline, then activate**—If you do not have Internet access, go to https://activate.doubletake.com from another machine that has Internet access. Complete and submit the activation form, and you will receive an e-mail with the activation key. Activate your server by entering the activation key in the **Add activation codes and activations keys** field and clicking **Add**.

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

If your activation codes needs to be activated, you will have 14 days to do so.

If you 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). See *Deactivating licenses* on page 58. 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.

- **Total licenses quantity**—If your server is a Double-Take controller appliance, you will see the available and total license count displayed. The controller appliance handles all license management for your agentless vSphere jobs. Each agentless vSphere job will use one license from your available quantity. If you create jobs for five virtual machines at one time, five licenses will be used. When you delete a job, the license associated with the job will be released to be available for another job. Stopping a job will not delete the license associated with the job. If you have no additional licenses available, you cannot create any new jobs. You can add to your available license quantity by purchasing additional licenses.

Server setup properties

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

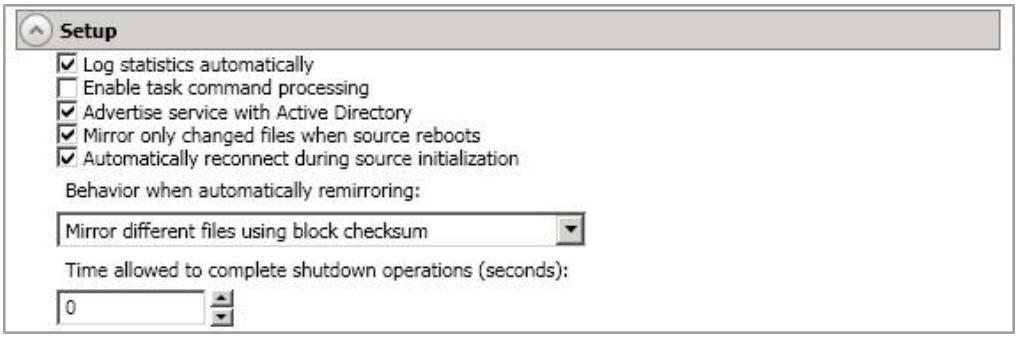

- **Log statistics automatically**—If enabled, Double-Take statistics logging will start automatically when Double-Take is started.
- **Enable task command processing**—Task command processing is a Double-Take feature that allows you to insert and run tasks at various points during the replication of 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 Windows NTFS change journal to track file changes. If the source is rebooted, only the files identified in the change journal will be remirrored to the target. This setting helps improve mirror times.
- **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.
  - **Do not mirror**—Do not automatically remirror any files. If you select this option, you will have to start a mirror manually to guarantee data integrity.
  - **Mirror different files using block checksum**—Any file that is different on the source and target based on date, time, and/or size is flagged as different. The mirror then performs a checksum comparison on the flagged files and only sends those blocks that are different.
  - **Mirror all files**—All files are sent to the target.
  - **Mirror different files**—Any file that is different on the source and target based on date, time, and/or size is sent to the target.
  - **Mirror only newer files**—Only those files that are newer on the source are sent to the target.

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

- **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.

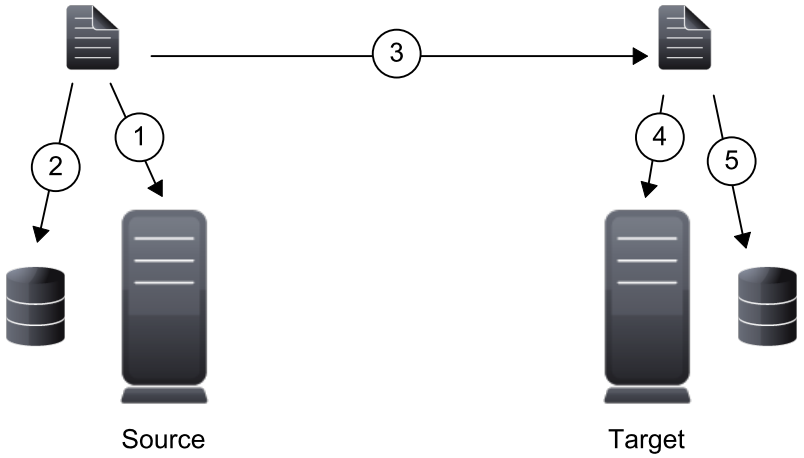

1. If data cannot immediately be transmitted to the target, it is stored in system memory. You can configure how much system memory you want Double-Take to use for all of its processing.
2. When the allocated amount of system memory is full, new changed data bypasses the full system memory and is queued 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.

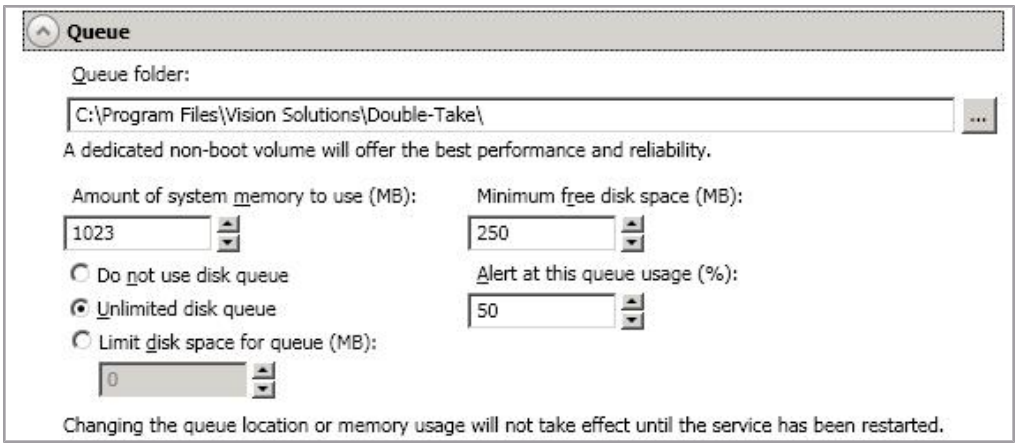

- **Queue folder**—This is the location where the disk queue will be stored. Any changes made to the queue location will not take effect until the Double-Take service has been restarted on the server.

  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.

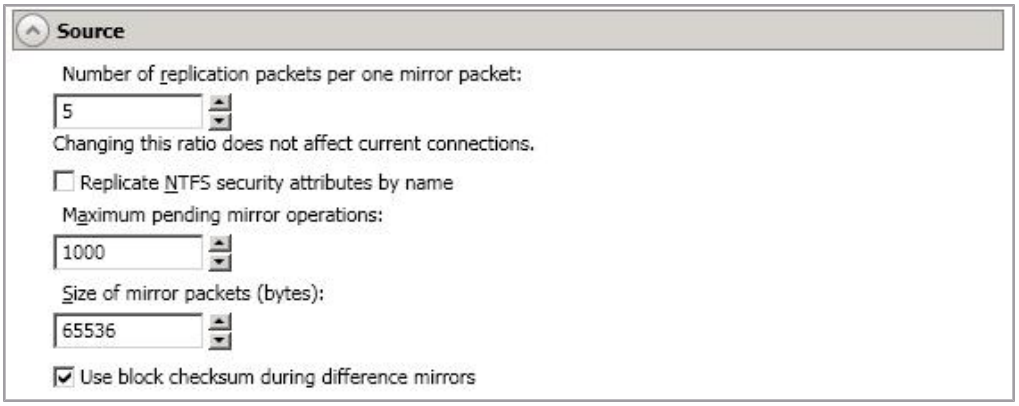

- **Number of replication packets per one mirror packet**—You can specify the ratio of replication packets to mirror packets that are placed in the source queue. The default value (5) allows Double-Take to dynamically change the ratio as needed based on the amount of replication data in queue. If you 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**—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. Additionally, you must enable this option on your target server before starting a restoration, because the target is acting like a source during a restoration.

  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.
- **Use block checksum during difference mirrors**—This option allows a file difference mirror to check each block of data, regardless of the file attributes. If this option is disabled, Double-Take will assume files are synchronized if their attributes match.

Database applications may update files without changing the date, time, or file size. Therefore, if you are using database applications, you should enable **Use block checksum during difference mirrors** to ensure proper file comparisons.

If you are not using database applications, disabling this option will shorten mirror times.

File Differences Mirror Options Compared

The following table will help you understand how the source server block checksum option works together with the various difference mirror job options. See the instructions for creating your job type to see which mirroring job options are available for that job type.

An X in the table indicates that option is enabled. An X enclosed in parentheses (X) indicates that the option can be on or off without impacting the action performed during the mirror.

Not all job types have the source newer option available.

Source Server Properties	Job Properties			Action Performed
Block Checksum Option	**File Differences Option**	**Source Newer Option**	**Block Checksum Option**	
(X)	X			Any file that is different on the source and target based on the date, time, size, and/or attribute is transmitted to the target. The mirror sends the entire file.
(X)	X	X		Any file that is newer on the source than on the target based on date and/or time is transmitted to the target. The mirror sends the entire file.
	X		X	Any file that is different on the source and target based on date, time, size, and/or attributed is flagged as different. The mirror then performs a checksum comparison on the flagged files and only sends those blocks that are different.
X	X		X	The mirror performs a checksum comparison on all files and only sends those blocks that are different.
(X)	X	X	X	Any file that is newer on the source than on the target based on date and/or time is flagged as different. The mirror then performs a checksum comparison on the flagged files and only sends those blocks that are different.

Target server properties

These properties are specific to the target server role.

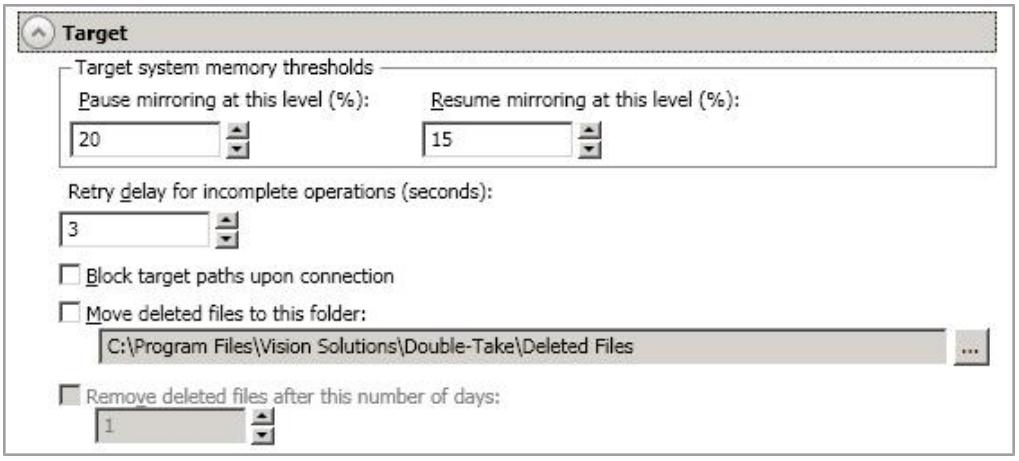

- **Pause mirroring at this level**—You can specify the maximum percentage of Windows system memory that can contain mirror data before the target signals the source to pause the sending of mirror operations. The default setting is 20.
- **Resume mirroring at this level**—You can specify the minimum 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.
- **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. After failover, any target paths that are blocked will be unblocked automatically during the failover process so that users can modify data on the target after failover. During restoration, the paths are automatically blocked again. If you failover and failback without performing a restoration, the target paths will remain unblocked.

  

  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. (SSL or TLS are not supported.) 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.

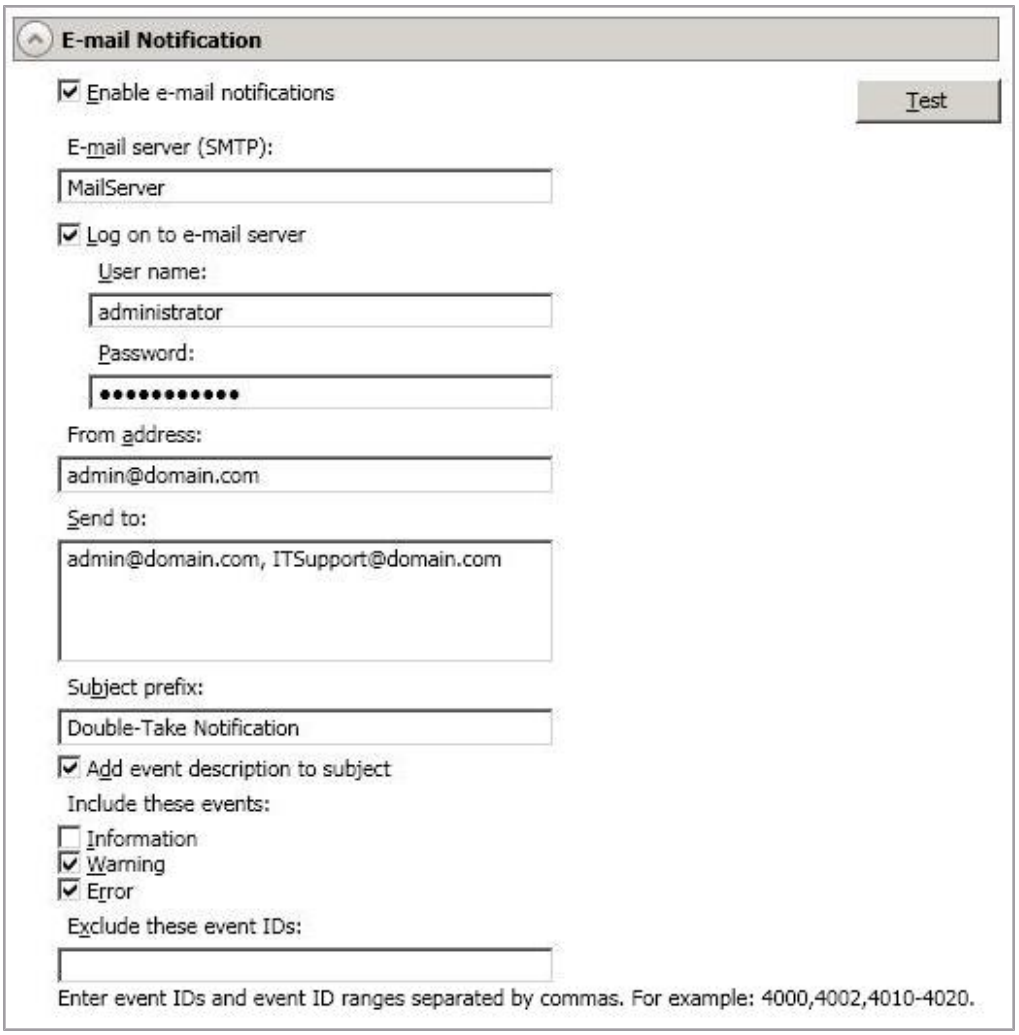

- **Enable e-mail notification**—This option enables the e-mail notification feature. Any specified notification settings will be retained if this option is disabled.
- **E-mail server**—Specify the name of your SMTP mail server.
- **Log on to e-mail server**—If your SMTP server requires authentication, 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 255 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 and failover.

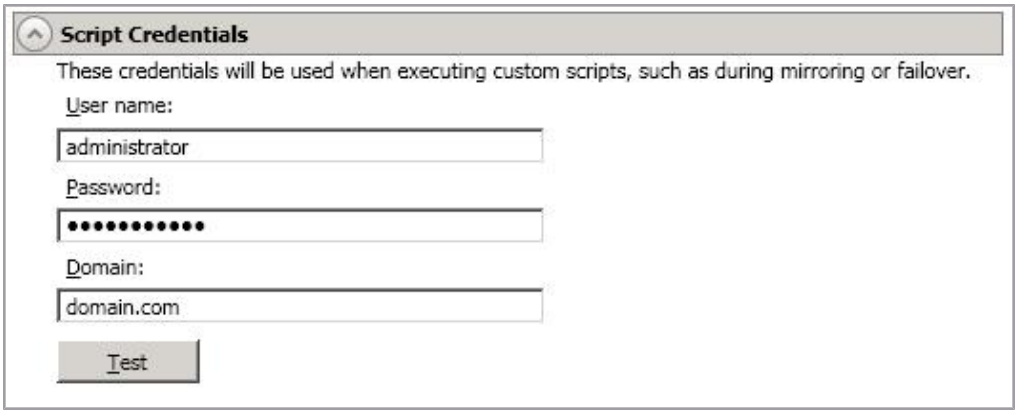

Specify a **User name**, **Password**, and **Domain** to use when running the scripts. If you do not specify any security credentials, the account running the Double-Take service will be used. After you have specified credentials, you can click **Test** to confirm the credentials can be used for a successful login. It 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.

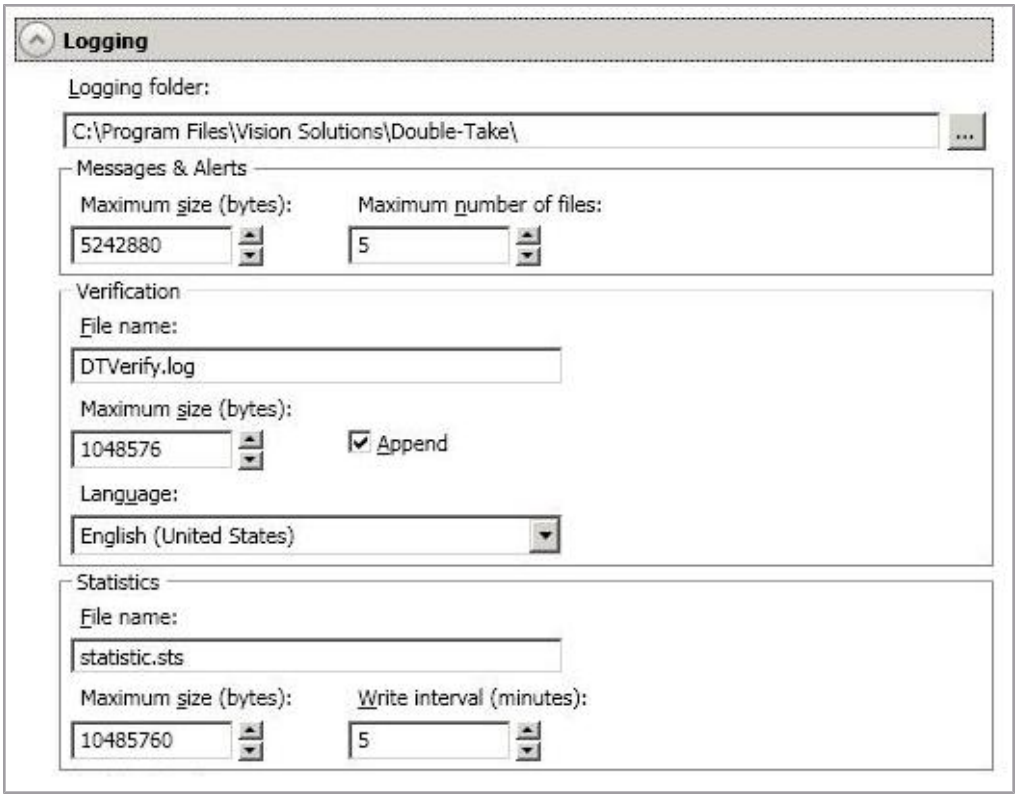

- **Logging folder**—Specify the directory where each of the log files in this section are stored. The default location is the directory where the Double-Take program files are installed.
- **Messages & Alerts**—These settings apply to the service log file.
  - **Maximum size**—Specify the maximum size, 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 5, 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 102.
  - **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**—At this time, English is the only language available.

- **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 *Statistics* on page 688.
  - **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: 1/24/2013 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 = 1/17/2013 8:21:36 PM Size = 1272 Mask = [0x20]
     Target Attributes: Timestamp = 1/17/2013 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 = 1/17/2013 8:21:37 PM Size = 1272 Mask = [0x20]
     Target Attributes: Timestamp = 1/17/2013 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 = 1/12/2013 3:28:20 PM Size = 17 Mask = [0x20]
     Target Attributes: Timestamp = 1/20/2013 5:05:26 PM Size = 2 Mask = [0x20]
      17 BYTES OUT OF SYNC
Completion Time: 1/24/2013 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.

Server and job settings

The easiest way to view and change select server and job settings is through the Double-Take Console. However, not all of the settings are available there. To view and update the remaining settings, in addition to the settings available in the console, you will need to go to HKEY_LOCAL_MACHINE\SOFTWARE\NSI Software\Double-Take\CurrentVersion in the registry. For a Linux server, you can use DTSetup to modify the configuration settings.

The following table lists all of the settings, in decimal value.

Double-Take Availability and Double-Take Move share the same set of server and job settings. Some settings apply to one product, some to the other, and some to both. For settings that apply to both, the Double-Take Availability terminology is used. For example, PreFailoverScript is used for the script to be run before failover or cutover.

If you are using a Linux source server, you will have only a subset of the settings listed below.

AcquireDataRetryLimit

Description—The length of time, in milliseconds, spent retrying a file read if there is a read error

Values—Any positive, integer value

Default—2000

Console Setting—None

Service restart required—No

ActivationCode

Description—24-character Double-Take activation code

Values—Unique value for each customer

Default—N/A

Console Setting—Edit Server Properties page, Licensing section, Current activation codes

Service restart required—No

AdapterFlags

Description—Specifies the adapter to use when establishing a connection. This option should not be changed.

Values—2 Encryption, 4 Network Data Representation

Default—4

Console Setting—None

Service restart required—Yes

AddOnCodes

Description—This setting is no longer used.

Advertisement

Description—Indicates if the server uses Active Directory to advertise itself so that the Double-Take Console can be populated through automatic discovery

Values—0 Do not use Active Directory advertisement, 8 Use Active Directory advertisement

Default—8

Console Setting—Edit Server Properties page, Setup section, Advertise service with Active Directory

Service restart required—Yes

Notes—If Active Directory advertisement is enabled, there is a 200 byte impact on the Active Directory service for each server that registers. The Double-Take service registers with Active Directory at startup and unregisters at shutdown.

AllFailover

Description—Specifies which IP addresses to failover

Values—0 Failover only monitored IP addresses, 1 Failover all IP addresses

Default—1

Console Setting—Set Options page, Failover Options section, Failover IP addresses

Service restart required—No

AllMustFail

Description—Specifies whether or not all IP addresses must fail for failover to take place

Values—0 any IP address can fail, 1 All IP addresses must fail

Default—1

Console Setting—None

Service restart required—No

ArchiveExclusionDirectories

Description—This setting is no longer used.

ArchiveExclusionFiles

Description—This setting is no longer used.

ArchiveLoopAttempts

Description—This setting is no longer used.

ArchiveLoopDelay

Description—This setting is no longer used.

ArchivePreviewFileName

Description—This setting is no longer used.

ArchivePreviewMaxFileSize

Description—This setting is no longer used.

ArchiveRemoveBinFileOnRecall

Description—This setting is no longer used.

ArchiveRequireMirrorCompleteTime

Description—This setting is no longer used.

ArchiveUseAlternateDate

Description—This setting is no longer used.

ArchiveUseDNSName

Description—This setting is no longer used.

AutoCalcEulaAccepted

Description—Used internally by Double-Take. Do not modify this entry.

AutoReconnect

Description—Specifies whether to reinstate the target connection(s) when the source machine is brought online after a source machine failure

Values—0 Do not reconnect, 1 Reconnect

Default—1

Console Setting—Edit Server Properties page, Setup section, Automatically reconnect during source initialization

Service restart required—Yes

AutoRemirror

Description—Specifies whether to remirror when a source is brought online after an auto-disconnect

Values—0 Do not remirror, 1 Perform a file differences checksum mirror, 2 Perform a full mirror, 3 Perform a file differences mirror, 4 Perform a date comparison mirror and send data only if the source data is newer than the target data.

Default—1

Console Setting—Edit Server Properties page, Setup section, Behavior when automatically reconnecting

Service restart required—No

AutoRemirrorRetry

Description—Specifies how often, in seconds, the source should check for connections that have been reconnected but still need to be remirrored

Values—any integer

Default—30

Console Setting—None

Service restart required—No

AutoRetransmit

Description—Determines whether or not a source that has lost its connection with a target will attempt to reconnect to the target

Values—0 Do not attempt to reconnect, 1 Attempt to reconnect

Default—1

Console Setting—None

Service restart required—No

BackupDir

Description—Location on the target of the backup of the protected data sets

Values—any valid path

Default—the location where the Double-Take files were installed

Console Setting—None

Service restart required—No

CalculateByVolume

Description—Calculates the approximate size of a protected data set by using the size of the volume and subtracting the free space

Values—0 Disabled, 1 Enabled

Default—0

Console Setting—None

Service restart required—Yes

Notes—If your protected data set contains a large number of files, for example, 250,000 or more, you may want to disable the calculation of the protected data set size so that data will start being mirrored sooner. If calculation is enabled, the source calculates the file size before it starts mirroring. This can take a significant amount of time depending on the number of files and system performance. Disabling calculation will result in the mirror status not showing the percentage complete or the number of

bytes remaining to be mirrored. CalculateByVolume can be enabled as a workaround. This setting will get the amount of disk space in use for the entire volume from the operating system, so the calculation occurs instantaneously. However, if the entire volume is not being replicated, the mirror percentage complete and bytes remaining will be incorrect accordingly.

Do not enable this option if you are using one of the following job types because it will bypass needed hard link processing: full server, full server to ESX, V to ESX, full server to Hyper-V, or V to Hyper-V.

CalculateOnConnect

Description—Specifies whether or not the amount of data to be mirrored should be calculated on connection

Values—0 Do not calculate on connection, 1 Calculate on connection

Default—1

Console Setting—None

Service restart required—Yes

CaseSensitiveRepSetQueries

Description—This entry is no longer used.

ChangeJournalState

Description—An internal setting for change journal tracking. Do not modify this setting.

ChangeJournalSystemState

Description—An internal setting for change journal tracking. Do not modify this setting.

ChecksumAll

Description—Setting to allow for the difference checksum option on mirror, verify, or restore to ignore the date, time, and size of the file and perform a checksum calculation on all files

Values—0 Checksum using date, time, size comparison, 1 Checksum all files regardless of the date, time, or file size

Default—1

Console Setting—Edit Server Properties page, Source section, Use block checksum during difference mirrors

Service restart required—No

ClusterDir

Description—Location of a Microsoft Cluster Service installation, if it exists

Values—any valid path

Default—determined by the Microsoft Cluster Service installation

Console Setting—None

Service restart required—No

ConnectionFile

Description—Name of the database file containing connection information

Values—any valid file name

Default—connect.sts

Console Setting—None

Service restart required—No

CreateDumpOnAckErrors

Description—Enables additional logging for out of order acknowledgement errors

Values—0 Do not create a logging file, 1 Create a logging file

Default—0

Console Setting—None

Service restart required—No

DataPath

Description—The location of the Double-Take file attribute, protected data set, connection, and schedule database files

Values—any valid path

Default—the location where the Double-Take files were installed

Console Setting—None

Service restart required—No

DefaultAddress

Description—The default primary IP address in a multi-homed server

Values—any valid IP address that will act as your primary IP address for connecting the source to the target

Default—<null>

Console Setting—Edit Server Properties page, General section, Default address

Service restart required—Yes

DefaultProtocol

Description—The default protocol

Values—2 IPv4 protocol only, 23 IPv4 and IPv6 protocols, 3 TDU (Throughput Diagnostics Utility)

Default—2 for Windows 2003, 23 for Windows 2008 and 2012

Console Setting—None

Service restart required—Yes

DefaultReaderType

Description—Internal setting used by Double-Take RecoverNow for recoveries. Do not modify this setting.

DelayGCArbitration

Description—Number of seconds to delay the arbitration process. This option allows time for the network to become stable before trying to execute arbitration logic, for example, when a cluster failover has occurred, but the network has a lag before it returns to a stable state. Arbitration should not start until the network is back in that stable state.

Values—any positive number

Default—0

Console Setting—None

Service restart required—No

DelayGCConnection

Description—Delays the GeoCluster Replicated Disk resource connection to allow the cluster service enough time to reset

Values—1-15

Default—3

Console Setting—None

Service restart required—No

DiffMirrorHardLinkCleanup

Description—Specifies if files with more than one hard link are deleted on the target during a difference mirror and then relinked after the remirror is complete. This setting only applies to Windows 2008 and 2012 servers with a full server job or full server migration job. If mirror performance is negatively impacted by this setting, you may want to disable it.

Values—0 Hard link files are not deleted and relinked during a difference mirror, 1 Hard link files are deleted and relinked during a difference mirror

Default—1

Console Setting—None

Service restart required—No

DirUNetPort

Description—Port used by pre-5.2 versions for directed UDP communications

Values—1025 - 65535

Default—1105

Console Setting—None

Service restart required—Yes

DisableAttributeReplication

Description—Specifies whether or not attributes (read-only, hidden, and so on) are replicated to the target

Values—0 Enable attribute replication, 1 Disable attribute replication

Default—0

Console Setting—None

Service restart required—No

DropOpOnAccessDeniedError

Description—Specifies whether or not operations are dropped or retried after an access denied error

Values—0 The operation will be retried, 1 The operation will be dropped

Default—1

Console Setting—None

Service restart required—No

DropOpOnHandleError

Description—Determines if an additional attempt is made to access a file by a Microsoft API call if the Double-Take call fails.

Values—0 When opening a file using the Double-Take driver fails, attempt to open the file usingthe Microsoft Win32 API, 1 When opening a file using the Double-Take driver fails, skip the file and document it in the Double-Take log

Default—1

Console Setting—None

Service restart required—No

Notes—If the value is set to 0 and the Win32 call also fails, Double-Take will skip the file and document it in the Double-Take log

DTSetupType

Description—Used by the Double-Take installation program to maintain the installation settings for an upgrade. Do not modify this setting.

DumpDiskQuotaIntervalMinutes

Description—Specifies how often, in minutes, a snapshot of the disk quota is taken as a backup in case the live registry is not usable at failover or cutover

Values—any integer

Default—240

Console Setting—None

Service restart required—No

DumpHiveIntervalMinutes

Description—Specifies how often, in minutes, a snapshot of the registry is taken as a backup in case the live registry is not usable at failover or cutover

Values—any integer

Default—240

Console Setting—None

Service restart required—No

EmailEnabled

Description—Specifies if e-mail notification is enabled

Values—0 E-mail notification is disabled, 1 E-mail notification is enabled

Default—0

Console Setting—Edit Server Properties page, E-mail Notification section, Enable e-mail notifications

Service restart required—Yes from the registry, No from the console

Notes—This is a read-only setting. If you change this setting using the registry editor, e-mail notification will not automatically start. You must use the console or a PowerShell script to start e-mail notification.

EmailExcludeIds

Description—Identifies the Windows Event Viewer messages that are excluded from e-mail notification.

Values—Comma or semicolon separated list of Event Viewer IDs. You can indicate ranges within the list.

Default—None

Console Setting—Edit Server Properties page, E-mail Notification section, Exclude these event IDs

Service restart required—Yes from the registry, No from the console

EmailFromAddress

Description—Specifies the e-mail address that will appear in the From field of Double-Take generated e-mail messages.

Values—Any valid e-mail address, up to 256 characters

Default—None

Console Setting—Edit Server Properties page, E-mail Notification section, From address

Service restart required—No

EmailIncludeCategories

Description—Specifies which Event Viewer messages are sent via e-mail

Values—1 Error messages will be sent via e-mail, 2 Warning messages will be sent via e-mail, 3 Information messages will be sent via e-mail

Default—1,2

Console Setting—Edit Server Properties page, E-mail Notification section, Include these events

Service restart required—Yes from the registry, No from the console

EmailNotificationList

Description—Specifies the e-mail address(es) that will receive Double-Take generated e-mail messages.

Values—A comma separated list of valid e-mail addresses, up to 256 addresses. Each address is limited to 256 characters.

Default—None

Console Setting—Edit Server Properties page, E-mail Notification section, Send to

Service restart required—No

EmailPassword

Description—The password required for SMTP server authentication

Values—Any valid password text

Default—None

Console Setting—Edit Server Properties page, E-mail Notification section, Password

Service restart required—No

Notes—Since the password is encrypted for security, this entry cannot be displayed or changed through the registry.

EmailServer

Description—The name of the SMTP server for e-mail notification

Values—Any valid server name text

Default—None

Console Setting—Edit Server Properties page, E-mail Notification section, E-mail server (SMTP)

Service restart required—No

EmailServerPort

Description—Specifies the port that the SMTP e-mail server is using

Values—any valid port number

Default—25

Console Setting—None

Service restart required—No

EmailSmtpLogin

Description—Specifies if SMTP server authentication for e-mail notification is enabled or disabled

Values—0 SMTP authentation is disabled, 1 SMTP authenticaion is enabled

Default—0

Console Setting—Edit Server Properties page, E-mail Notification section, Log on to e-mail server

Service restart required—No

Notes—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.

EmailSubjectDesc

Description—Specifies if the Event Viewer message will be appended to the end of the subject line for e-mail notification

Values—0 Event Viewer message is not included in the subject line, 1 Event Viewer message is included in the subject line

Default—1

Console Setting—Edit Server Properties page, E-mail Notification section, Add event description to subject

Service restart required—No

EmailSubjectPrefix

Description—Specifies unique text which will be inserted at the front of the subject line for each Double-Take generated e-mail message. This will help distinguish the Double-Take messages from other messages.

Values—Any valid text

Default—Double-Take Notification

Console Setting—Edit Server Properties page, E-mail Notification section, Subject prefix

Service restart required—No

EmailUsername

Description—The user ID required for SMTP server authentication

Values—Any valid user ID text

Default—None

Console Setting—Edit Server Properties page, E-mail Notification section, User name

Service restart required—No

Notes—Since the username is encrypted for security, this entry cannot be displayed or changed through the registry.

EnableCRCCheck

Description—Indicates if Double-Take will perform a cyclic redundancy check between the source and target to identify corrupted packets

Values—0 Disabled, 1 Enabled

Default—0

Console Setting—None

Service restart required—No

Notes—This option only needs to be set on the source server. However, if you will be restoring or reversing, where the roles of the servers are reversed, then you will need to set this option on the target as well.

EnableDHCP

Description—Indicates if Double-Take DHCP support is enabled

Values—0 Disabled, 1 Enabled

Default—1

Console Setting—None

Service restart required—No

EnableEFSVerify

Description—Indicates if Double-Take will verify Microsoft encryption on the source before transmitting the encrypted file to the target

Values—0 Disabled, 1 Enabled

Default—0

Console Setting—None

Service restart required—No

EnableFileOpenTracing

Description—Specifies if debug-level messages are enabled to trace all mirroring and replicated files that are opened

Values—0 Do not trace files that are opened, 1 Trace files that are opened

Default—0

Console Setting—None

Service restart required—Yes

Notes—This option should only be enabled (1) for temporary, debug sessions as instructed by technical support.

EnablePerformanceTracking

Description—This entry will be used in the future.

EnableRootEncryption

Description—Specifies if the top-level folders of a protected data set are encrypted on the source, they will be encrypted on the target as well

Values—0 Disabled, 1 Enabled

Default—1

Console Setting—None

Service restart required—No

Notes—If the top-level folders in a protected data set are not encrypted, disabling this option may obtain a small performance improvement.

EnableShortFileNameProcessing

Description—Indicates if Double-Take will correct any short file names created by the operating system on the target during a mirror or for create and rename operations

during replication

Values—0 Do not correct any short file names on the target, 1 Correct short file names on the target

Default—0

Console Setting—None

Service restart required—No

EnableSnapshots

Description—Specifies whether Double-Take snapshot functionality is enabled

Values—0 Double-Take snapshot functionality is disabled, 1 Double-Take snapshot functionality is enabled

Default—1

Console Setting—None

Service restart required—Yes

Notes—This setting only impacts Double-Take snapshot functionality. If this setting is disabled, other snapshot software such as Microsoft Volume Shadow Copy will be not be impacted.

EnableTaskCmdProcessing

Description—Queues tasks inline with replication data

Values—0 Disable task command processing, 1 Enable task command processing

Default—0

Console Setting—Edit Server Properties page, Setup section, Enable task command processing

Service restart required—No

EncryptNetworkData

Description—Encrypts Double-Take data before it is sent from the source to the target

Values—0 Disable data encryption, 1 Enable data encryption

Default—0

Console Setting—None

Service restart required—No

Notes—Both the source and target must be Double-Take encryption capable (Double-Take 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.

ExtensionNumber

Description—Used by the Double-Take log files.

FailbackHostname

Description—Returns the host SPN (Service Principle Name) to its original setting on failback

Values—0 Disabled, 1 Enabled

Default—0

Console Setting—None

Service restart required—No

Notes—If you are using Active Directory, this option should be enabled or you may experience problems with failback.

FailoverData1

Description—An internal setting for failover. Do not modify this setting.

FailoverData2

Description—An internal setting for failover. Do not modify this setting.

FailoverHostname

Description—Automatically removes the host SPN (Service Principle Name) from Active Directory on the source

Values—0 Disabled, 1 Enabled

Default—0

Console Setting—None

Service restart required—No

Notes—If you are using Active Directory, this option should be enabled or you may experience problems with failover.

FailoverOnRouteFailure

Description—Determines if failover will occur when receiving a router message back from an IP address on the network

Values—0 Failover will not occur when receiving a destination host unreachable message, 1 Failover will occur when receiving a destination host unreachable message

Default—1

Console Setting—None

Service restart required—No

FCCHelpPath

Description—This entry is no longer used.

FileAccessRetry

Description—The number of times a failed driver call will be retried by the service.

Values—1 - 65535

Default—10

Console Setting—None

Service restart required—No

FileQueueSize

Description—When a mirror is started, one thread reads from the disk and builds the file queue. Another set of threads reads files off of the queue and sends them to the target. This setting is the maximum size of the queue in entries. If you had 100 files to be mirrored and this was set to 16 (the default value), the first thread would fill the queue to a maximum of 16 entries.

Values—1 - 65535

Default—16

Console Setting—None

Service restart required—No

Notes—This value must be set prior to starting the mirror process. The higher the number, the more memory that is used.

ForceReplaceOnFailover

Description—Specifies additional failover options

Values—0 Use standard failover add / replace settings with no additional settings, 1 Replace the target server name with that of the source and add the source IP address, 2 Add the source server name to the target and replace the target IP address, 3 Replace the target server name with that of the source and replace the target IP address

Default—0

Console Setting—None

Service restart required—No

ForceVerifyOnMirror

Description—Specifies if verification will be performed with every difference mirror

Values—0 Verification is not performed with every difference mirror, 1 Verification is performed with every difference mirror

Default—0

Console Setting—None

Service restart required—No

GenerateDumpOnShutdownCrashes

Description—Specifies if a log file will be created during a shutdown crash

Values—0 No log file is created and the default exception handler is used, 1 Log file is created

Default—0

Console Setting—None

Service restart required—Yes

HardLinkInterval

Description—Specifies the length of time, in seconds, to generate a hard link report

Values—any valid integer

Default—3600

Console Setting—None

Service restart required—No

HardLinkLogPath

Description—Specifies the location where hard links will be logged. If no path is specified, the location defined in LogDir will be used.

Values—any valid path

Default—None

Console Setting—None

Service restart required—No

HBLoopback

Description—This entry is no longer used.

HBTrace

Description—Specifies whether heartbeat debugging information is generated

Values—0 not generated, 1 Generated

Default—0

Console Setting—None

Service restart required—No

HBTTL

Description—Number of seconds without receiving a heartbeat before a remote machine is considered unavailable

Values—0 - 65535

Default—10

Console Setting—None

Service restart required—No

HeartbeatIgnoreIPs

Description—This setting is no longer used.

HPQueueRatio

Description—Ratio of replication packets to one mirror packet

Values—1 - 65535

Default—5

Console Setting—Edit Server Properties page, Source section, Number of replication packets per one mirror packet

Service restart required—No for future connections, Yes for the current connection

Notes—An HPQueueRatio of 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.

IgnoreAlternateStreamFiles

Description—Specifies alternate streams to skip during mirroring and replication

Values—a semi-colon separate list of stream names. The stream names are not case-sensitive

Default—none

Console Setting—None

Service restart required—No

IgnoreArchiveBit

Description—Specifies if the archive bit is compared during verification

Values—0 Archive bit is compared during a verification, 1 Archive bit is not compared during a verification

Default—1

Console Setting—None

Service restart required—No

IgnoreDeleteOps

Description—Specifies if file and directory delete operations will be replicated to the target

Values—0 Delete operations are replicated to the target, 1 Delete operations are not replicated to the target

Default—0

Console Setting—None

Service restart required—No

IgnoreOpLockErrors

Description—Specifies how files that are locked open on the source are handled during mirroring

Values—0 Fail the mirror and record OpLock errors in the log. The job state will be set to mirror required, 1 Ignore the lock errors and continue the mirror. This option does not guarantee data integrity. There may be differences in the file that was locked.

Default—0

Console Setting—None

Service restart required—No

IgnorePPPAddresses

Description—Identifies if Double-Take will use PPP (Point-to-Point Protocol) or SLIP (Serial Line Internet Protocol) adapters

Values—0 Double-Take will send out heartbeats across the PPP/SLIP adapter, 1 Double-Take will not send out heartbeats across the PPP/SLIP adapter

Default—1

Console Setting—None

Service restart required—No

IgnoreSourceErrors

Description—Ignores source errors that will cause an update to the target data state

Values—0 Do not ignore source errors, 1 Ignore source errors

Default—0

Console Setting—None

Service restart required—No

IgnoreThumbnailStreams

Description—Specifies if thumbnails will be replicated to the target.

Values—0 Double-Take will mirror and replicate all data streams, 1 Double-Take will not mirror or replicate any data about the alternate data streams for thumbnail images. When comparing data for a verification or difference mirror,alternate data streams for thumbnails will not be reported as different.

Default—1

Console Setting—None

Service restart required—If you change this value to 0, you must restart the Double-Take service in order for the Double-Take driver to begin sending all data stream information to the service. If you change this value to 1, you do not need to restart the service.

IgnoreWriteFailureOnTarget

Description—Specifies whether failures to write a file on the target are logged

Values—0 Log all write failures on the target, 1 or any larger integer indicates that number of write failures which will be ignored before starting to log the write failures

Default—0

Console Setting—None

Service restart required—No

IncludeSysVolInfo

Description—Specifies whether the system volume information folder is mirrored and replicated

Values—0 Do not include the system volume information folder, 1 Include the system volume information folder

Default—0

Console Setting—None

Service restart required—No

InstallPath

Description—Path specified during the Double-Take installation. Do not modify this entry.

InstallVersionInfo

Description—Installation number specified during the Double-Take installation. Do not modify this entry.

InteractWithDesktopForFOScripts

Description—Runs the failover scripts interactively on the user's local session

Values—0 Do not run the scripts interactively, 1 Run the scripts interactively.

Default—0

Console Setting—None

Service restart required—Yes

IntermediateQueueLimit

Description—Amount of memory, in KB, that may be allocated to the intermediate queue by the system memory manager when MemoryAllocatorMode is set to mixed mode (2).

Values—512-4194304

Default—65536

Console Setting—None

Service restart required—Yes

IPFailover

Description—Specifies whether or not to failover the IP addresses during failover

Values—0 Do not failover IP addresses 1 Failover IP addresses

Default—1

Console Setting—Set Options page, Failover Options section, Failover IP addresses

Service restart required—No

KFAIOpenRetry

Description—Specifies the number of times an operation is retried if the driver return an error

Values—any valid integer

Default—10

Console Setting—None

Service restart required—No

LanguagesAvailable

Description—Specifies the Double-Take language support that has been installed. Do not modify this setting. If you need to add or remove language support, use the Double-Take installation program.

LanguageSelected

Description—Specifies the language of the verification log

Values—Depends on LanguagesSupported

Default—Language used during the installation

Console Setting—Edit Server Properties page, Logging section, Language

Service restart required—Yes

LanguagesSupported

Description—Specifies the available languages for the verification log. Currently English is the only language available. Do not modify this setting.

LastModifiedReadDelay

Description—Specifies the length of time, in seconds, to wait before reading the last modified file time attribute

Values—any valid integer

Default—15

Console Setting—None

Service restart required—No

Notes—This option is only used if SendLastModifiedTimeOnClose is disabled

LoadSourceTarget

Description—Specifies the functionality of the Double-Take service

Values—0 Neither the source nor target modules are loaded, 1 Only the source module is loaded, 2 Only the target module is loaded, 3 Both the source and target modules are loaded

Default—3

Console Setting—None

Service restart required—Yes

LogAllOrphans

Description—This entry is no longer used.

LogDir

Description—The location of the Double-Take messages/alerts, verification, and statistics log files

Values—any valid path

Default—the location where the Double-Take files were installed

Console Setting—Edit Server Properties page, Logging section, Logging folder

Service restart required—Yes

LogFile

Description—The name of the Double-Take messages/alerts log file

Values—any valid file name

Default—DTLog

Console Setting—None

Service restart required—No

LogHardlinks

Description—Indicates whether hard links are logged to replication_set_name.log when the protected data set size is calculated

Values—0 Hard links are not logged, 1 Hard links are logged

Default—0

Console Setting—None

Service restart required—No

LogMessageLevel

Description—Specifies the types of messages logged to the.dtl files

Values—0 No messages will be logged, 1 Only alert messages will be logged, 2 Alert and release messages will be logged, 3 Alert, release, and debug messages will be logged

Default—2

Console Setting—None

Service restart required—No

LoopbackID

Description—This entry is no longer used.

MaxChecksumBlocks

Description—Specifies the number of checksum values retrieved from the target

Values—any integer

Default—32

Console Setting—None

Service restart required—No

MaxConnections

Description—Number of network requests that can be processed simultaneously. Windows is limited to 5 simultaneous requests.

Values—0 - 65535

Default—5

Console Setting—None

Service restart required—Yes

Notes—Vision Solutions recommends that you not change this value.

MaxLogFileSize

Description—Maximum size, in bytes, of any .dtl log file

Values—limited by available disk space

Default—5242880

Console Setting—Edit Server Properties page, Logging section, Maximum size (under Messages & Alerts)

Service restart required—No

MaxLogPathname

Description—The maximum length of a file name (the entire volume\directory\filename including slashes, spaces, periods, extensions, and so on) that will be displayed in the Double-Take log file and the Windows Event Viewer. File names longer than the MaxDisplayablePath will be truncated and will be followed by an ellipsis (...).

Values—1-32760

Default—32760

Console Setting—None

Service restart required—No

MaxNumberofLogFiles

Description—Maximum number of .dtl log files that can exist at one time. When Double-Take creates a new .dtl file, if this number is exceeded, the oldest .dtl file is deleted.

Values—1 - 999

Default—20

Console Setting—Edit Server Properties page, Logging section, Maximum number of files

Service restart required—No

MaxOpBufferSize

Description—An internal setting for memory buffering. Do not modify this setting.

MaxRemoveOrphansOpSize

Description—Determines whether or not Double-Take will send over multiple orphan operations. Double-Take will send over the operations if a directory has more files than this number.

Values—0 - 131072

Default—1000

Console Setting—None

Service restart required—No

MaxRetry

Description—A generic, application wide setting specifying the number of retry attempts for processes such as creating sockets, starting the service, and so on

Values—any integer

Default—5

Console Setting—None

Service restart required—Yes

MaxWriteChunkSize

Description—Maximum merged op size (in bytes) used during replication

Values—1 - 131072

Default—65536

Console Setting—None

Service restart required—No

MCHelpPath

Description—This entry is no longer used.

MemoryAllocatorCallbackMode

Description—Determines what action is taken when the MemoryQueueToDiskThreshold is met

Values—0 Auto-disconnect processing is initiated when theMemoryQueueToDiskThreshold has been met. Connections will be reestablished when auto-reconnect occurs, 1 The Double-Take service stops pulling operations from the driver when theMemoryQueueToDiskThreshold has been met. The target will pause the source.The service will resume pulling operations when the target tells the source to resume, 2 The source and target begin queuing operations to disk.

Default—2

Console Setting—None

Service restart required—Yes

MemoryQueueToDiskThreshold

Description—A percentage of QmemoryBufferMax that will trigger queuing to disk.

Values—any valid percentage

Default—75

Console Setting—None

Service restart required—Yes

MinCompressionFileSize

Description—The minimum file size, in bytes, that will be compressed. Files smaller than this size will not be compressed.

Values—any file size

Default—1024

Console Setting—None

Service restart required—No

MirrorChunkSize

Description—Block size, in bytes, used in the mirroring process

Values—1 - 1048576

Default—65536

Console Setting—Edit Server Properties page, Source section, Size of mirror packets

Service restart required—No

Notes—A higher block size value gives you better throughput, but only to a certain point, then it starts using more memory (this has to do with the way memory is allocated and deallocated). A lower block size value produces slower throughput, but uses memory efficiently.

MirrorEncryptedFiles

Description—Specifies if Windows 200x encrypted files are mirrored

Values—0 Encrypted files are not mirrored, 1 Encrypted files are mirrored

Default—1

Console Setting—None

Service restart required—No

MirrorOverwrite

Description—Determines if the mirror process overwrites existing files

Values—0 never overwrite, 1 always overwrite, 2 overwrite if older

Default—1

Console Setting—None

Service restart required—No

MirrorPrompting

Description—This entry is no longer used.

MirrorQueueLimit

Description—Maximum number of mirror operations that can be queued on the source machine

Values—1 - 65535

Default—1000

Console Setting—Edit Server Properties page, Source section, Maximum pending mirror operations

Service restart required—No

MirrorRootAttributes

Description—Specifies whether or not root permissions from the source are mirrored to the target

Values—0 Root permissions are not mirrored to the target, 1 Root permissions are mirrored to the target

Default—1

Console Setting—None

Service restart required—No

MirrorZeroKFiles

Description—Specifies whether or not empty files, zero byte files, are included in a mirror

Values—0 Zero byte files are skipped and not mirrored to the target, 1 All files are mirrored to the target

Default—1

Console Setting—None

Service restart required—No

Notes—If MirrorZeroKFiles is enabled (0), zero byte files are skipped during a full mirror, file differences mirror, and a verification with synchronization. Zero byte files that contain alternate data streams that are not empty, will still be skipped if MirrorZeroKFiles is enabled.

MissedHeartbeats

Description—Specifies the number of heartbeats that can go unanswered by the source before failover occurs, when using Double-Take service responses to monitor for failover

Values—1 - 65535

Default—20

Console Setting—None

Service restart required—No

Notes—This value is used in conjunction with the PingFrequency value to determine the value of Consider the source server failed (on the Set Options page, Failover Monitor section).

MissedPackets

Description—Specifies the number of requests sent by the target that go unanswered by the source before failover occurs, when using network responses to monitor for failover

Values—1 - 65535

Default—5

Console Setting—None

Service restart required—No

Notes—This value is used in conjunction with the PingFrequency value to determine the value of Consider the source server failed (on the Set Options page, Failover Monitor section).

MoveOrphanedFiles

Description—This entry is no longer used.

MoveOrphansDir

Description—This entry is no longer used.

NameFailover

Description—Specifies whether or not to failover machine names

Values—0 Do not failover machine names, 1 Failover machine names

Default—1

Console Setting—Set Options page, Failover Options section, Failover server name

Service restart required—No

NetPort

Description—Port used by pre-5.2 versions for TCP communications

Values—1025 - 65535

Default—1100

Console Setting—None

Service restart required—Yes

NetworkRetry

Description—Specifies the interval, in seconds, at which Double-Take will attempt to reconnect to the target

Values—any positive number

Default—10

Console Setting—None

Service restart required—No

NetworkStatusInterval

Description—An internal setting for network communications. Do not modify this setting.

NetworkTimeout

Description—The maximum length of time, in seconds, to wait on a network connection. If data is not received over a network connection within the specified time limit, the connection is closed. During idle periods, Double-Take sends small amounts of keep-alive data at an interval 1/6 of the NetworkTimeout value to keep the socket from being inadvertently closed.

Values—any integer

Default—120

Console Setting—None

Service restart required—No

Notes—If you are archiving files and it takes longer than the NetworkTimeout specified (for example, this may happen if the DTArchiveBin is located on an alternate volume), the archive operation will complete on the target, but the full file will not be changed to a link on the source because the source detected the network timeout.

NodeLockedLicenseKey

Description—An internal setting for licensing. Do not modify this setting.

NodeLockedServerInfo

Description—An internal setting for licensing. Do not modify this setting.

OpBufferMax

Description—Specifies the number of operations that can be stored in the memory queue prior to queuing to disk

Values—0 There is no limit to the number of operations that can be stored in the memory queue, 1 or any larger integer

Default—200000

Console Setting—None

Service restart required—No

OpBuffersCount

Description—An internal setting for memory buffering. Do not modify this setting.

OpLogging

Description—Specifies whether operations from the Double-Take driver are logged

Values—0 Do not log operations, 1 Log operations

Default—0

Console Setting—None

Service restart required—Yes

OutOfOrderDiff

Description—The maximum number of operations that can be out of order before the connection is paused

Values—any integer

Default—10

Console Setting—None

Service restart required—No

Notes—The larger the value, the more memory the Double-Take service on the targe service will use.

PingFrequency

Description—Specifies, in seconds, how often a ping is sent to the source from a monitoring target

Values—1 - 65535

Default—5

Console Setting—None

Service restart required—No

Notes—This value is used in conjunction with the MissedHeartbeats or MissedPackets value to determine the value of Consider the source server failed (on the Set Options page, Failover Monitor section).

Port

Description—Port connection for core Double-Take communications

Values—1025 - 65535

Default—6320

Console Setting—Edit Server Properties page, General section, Port

Service restart required—Yes

PostFailbackScript

Description—Location on the target where the target post-failback script is located

Values—Any valid path

Default—<null>

Console Setting—Set Options page, Failover Options section, Script file (under Post-failback script)

Service restart required—No

PostFailbackScriptArgs

Description—Arguments to be used with the target post-failback script

Values—Any valid argument

Default—<null>

Console Setting—Set Options page, Failover Options section, Arguments (under Post-failback script)

Service restart required—No

PostFailoverScript

Description—Location on the target where the target post-failover script is located

Values—Any valid path

Default—<null>

Console Setting—Set Options page, Failover Options section, Script file (under Post-failover script)

Service restart required—No

PostFailoverScriptArgs

Description—Arguments to be used with the target post-failover script

Values—Any valid argument

Default—<null>

Console Setting—Set Options page, Failover Options section, Arguments (under Post-failback script)

Service restart required—No

PreFailbackScript

Description—Location on the target where the target pre-failback script is located

Values—Any valid path

Default—<null>

Console Setting—Set Options page, Failover Options section, Script file (under Pre-failback script)

Service restart required—No

PreFailbackScriptArgs

Description—Arguments to be used with the target pre-failback script

Values—Any valid argument

Default—<null>

Console Setting—Set Options page, Failover Options section, Arguments (under Pre-failback script)

Service restart required—No

PreFailbackWait

Description—Specifies whether or not to wait for the target pre-failback script to complete before finishing a failback

Values—0 Do not wait, 1 Wait

Default—0

Console Setting—Set Options page, Failover Options section, Delay until script completes (under Pre-failback script)

Service restart required—No

PreFailoverScript

Description—Location on the target where the target pre-failover script is located

Values—Any valid path

Default—<null>

Console Setting—Set Options page, Failover Options section, Script file (under Pre-failover script)

Service restart required—No

PreFailoverScriptArgs

Description—Arguments to be used with the target pre-failover script

Values—Any valid argument

Default—<null>

Console Setting—Set Options page, Failover Options section, Arguments (under Pre-failover script)

Service restart required—No

PreFailoverWait

Description—Specifies whether or not to wait for the target pre-failover script to complete before finishing a failover

Values—0 Do not wait, 1 Wait

Default—0

Console Setting—Set Options page, Failover Options section, Delay until script completes (under Pre-failover script)

Service restart required—No

ProductCode

Description—Used by the Double-Take installation program to maintain the installation settings for an upgrade. Do not modify this entry.

ProductName

Description—Used by the Double-Take installation program to maintain the installation settings for an upgrade. Do not modify this entry.

QJournalDir

Description—The location where the queue is stored.

Values—any valid path

Default—the location specified during the installation

Console Setting—Edit Server Properties page, Queue section, Queue folder

Service restart required—No

Notes—For best results and reliability, you should select a dedicated, non-boot volume. The queue should be stored on a fixed, local NTFS volume. This location also stores the Double-Take driver pagefile.

QJournalFileSize

Description—The size, in MB, of each queuing transaction log file.

Values—any valid file size, up to 4095 MB

Default—5

Console Setting—None

Service restart required—No

QJournalFreeSpaceMin

Description—The minimum amount of disk space, in MB, in the specified QJournalDir that must be available at all times.

Values—dependent on the amount of physical disk space available

Default—250

Console Setting—Edit Server Properties page, Queue section, Minimum free disk space

Service restart required—No

Notes—The QJournalFreeSpaceMin should be less than the amount of physical disk space minus QJournalSpaceMax.

QJournalPreload

Description—The number of operations being pulled from the disk queue at one time. Do not modify this setting.

QJournalSpaceMax

Description—The maximum amount of disk space, in MB, in the specified QJournalDir that can be used for Double-Take queuing. When this limit is reached, Double-Take will automatically begin the auto-disconnect process.

Values—dependent on the amount of physical disk space available

Default—Unlimited

Console Setting—Edit Server Properties page, Queue section, Limit disk space for queue

Service restart required—No

Notes—The unlimited setting allows the disk queue usage to automatically expand whenever the available disk space expands. Setting this option to zero (0) disables disk queuing. Even if you are using the unlimited option, Double-Take will only store 16,384 log files. If you are using the default 5MB file size, this is approximately 80GB of data. If you anticipate needing to be able to queue more data than this, you should increase the size of the log files.

QLogWriteThrough

Description—Specifies if the disk queues are write-through mode

Values—0 Disk queues are not write-through mode, 1 Disk queues are write-through mode

Default—0

Console Setting—None

Service restart required—No

Notes—While write-through mode may decrease the frequency of auto-disconnects, it may also decrease the performance of the source server.

QMemoryBufferMax

Description—The amount of Windows system memory, in MB, that, when exceeded, will trigger queuing to disk.

Values—minimum 512, maximum is dependent on the server hardware and operating system

Default—1024

Console Setting—Edit Server Properties page, Queue section, Amount of system memory to use

Service restart required—Yes

QueryOnQuorumFile

Description—Identifies if the Double-Take service will reopen closed files on the quorum drive

Values—0 The Double-Take service will not attempt to reopen a closed file on the quroum drive to get security descriptors or last modified times, 1 The Double-Take service will attempt to reopen a closed file on the quroum drive to get security descriptors or last modified times.

Default—1

Console Setting—None

Service restart required—No

QueueSizeAlertThreshold

Description—The percentage of the queue that must be in use to trigger an alert message in the Windows Event Viewer.

Values—any valid percentage

Default—50

Console Setting—Edit Server Properties page, Queue section, Alert at this queue usage

Service restart required—Yes

Registered

Description—This entry is no longer used.

RemoveAllOrphans

Description—This entry is no longer used.

RemoveOrphansTime

Description—This entry is no longer used.

RemoveSharesOnDisconnect

Description—Specifies if shares are removed on the target machine when a Double-Take protected data set is disconnected from a target or a source machine is manually shutdown by the administrator. (Shares are not removed if either the source or target machines fail.)

Values—0 Remove shares from the target, 1 Do not remove shares from the target

Default—1

Console Setting—None

Service restart required—No

ReplaceTarget

Description—Specifies whether or not to replace the target identity with the source identity during a failover

Values—0 Do not replace, 1 Replace

Default—0

Console Setting—None

Service restart required—No

ReplicateNtSecurityByName

Description—Determines whether or not Double-Take replicates permissions and attributes assigned to local (non-domain) users and groups

Values—0 Do not replicate by name, 1 Replicate by name

Default—0

Console Setting—Edit Server Properties page, Source section, Replicate NTFS security attributes by name

Service restart required—No

ReplicationDiskCheckScript

Description—Specifies the script to run if validation of the replication drive fails

Values—Any valid path and script file

Default—<null>

Console Setting—None

Service restart required—No

ReplicationDiskCheckTimeOut

Description—Specifies the interval, in seconds, between validation checks when ReplicationDiskCheckSript is populated

Values—any integer

Default—300

GUI Setting—None

Service restart required—No

RepSetDBName

Description—Name of the database that contains protected data set information

Values—any valid file name

Default—DblTake.db

Console Setting—None

Service restart required—No

RestoreOverwrite

Description—Determines if the restoration process overwrites existing files

Values—0 never overwrite, 1 always overwrite, 2 overwrite if older

Default—2

Console Setting—None

Service restart required—No

RestorePrompting

Description—This entry is no longer used.

RetentionFlag

Description—This entry will be used in the future.

RunDTInfoOnCutover

Description—Specifies if DTInfo is launched before a failover or cutover when protecting an entire server

Values—0 Do not launch DTInfo, 1 Launch DTInfo

Default—1

Console Setting—None

Service restart required—No

RunScriptatSnaptime

Description—If a script is specified, the script is launched on the target before Double-Take executes any snapshots. The snapshot will not be executed until the script has completed. If the script returns an error, the snapshot will still execute.

Values—any valid path and script name

Default—<null>

Console Setting—None

Service restart required—No

RunScriptInsteadofSnap

Description—Specifies if a script specified in RunScriptAtSnaptime is executed

Values—0 Execute script specified in RunScriptAtSnaptime, 1 Do not execute script specified in RunScriptAtSnaptime

Default—1

Console Setting—None

Service restart required—No

SaveStatFile

Description—Determines if the statistic.sts (statistics logging) file is saved or ovewritten

Values—0 overwrite, 1 saved as statistic-old.sts

Default—1

Console Setting—None

Service restart required—No

ScheduleFile

Description—Name of the database file that contains transmission scheduling information

Values—any valid file name

Default—Schedule.sts

Console Setting—None

Service restart required—Yes

ScheduleInterval

Description—The number of seconds to wait before checking the transmission schedules to see if transmission should be started or stopped

Values—1 - 3600

Default—1

Console Setting—None

Service restart required—Yes

SendDirLastModifiedTime

Description—Specifies if the last modified time for directories will be transmitted to the target during a difference mirror

Values—0 last modified time on directories will not be sent to the target, 1 last modified time on directories will be sent to the target

Default—1

Console Setting—None

Service restart required—No

SendFileTimesOnCreate

Description—Specifies whether a file is accessed twice so that the file's creation time can be modified to match the source

Values—0 The Double-Take service will not access newly created files that have not been modified. These files on the target will have the date and time of when the file was created on the target, 1 The Double-Take service will access newly created files. These files on the target will have the same date and time as the source.

Default—0

Console Setting—None

Service restart required—No

Notes—New files created on the source that have not been modified will have the date and time of when the file is created on the target. The date and time will be corrected to match the source's true file attributes when a remirror or verification modifies them to match the source or the file is modified by a user or application on the source. For example, if the source machine's clock is set to 2:00 PM and the target machine is set to 4:00 PM, a newly created file that has not been modified will have a time stamp of 4:00 PM when it is applied to the target. If this option is enabled (1), Double-Take will access the file twice, to correctly set the time to 2:00 PM to reflect the file's true attributes. If this option is disabled (0), Double-Take will not access the file twice, and the file will have the target time of 4:00 PM until it is modified (remirror, verification, or user or application update).

SendLastModifiedTimeOnClose

Description—Specifies that the last modified time attribute is sent when a file is closed

Values—0 Last modified time is sent when Double-Take has not received any additional operations for the file in the time period specified by LastModifiedReadDelay, 1 Last modified time is sent when a file is closed, which may not be immediately depending on system processing

Default—1

Console Setting—None

Service restart required—No

Notes—If system processing delays (such as the system cache manager not flushing quickly enough) are causing delays in processing the last modified time, you may want to consider disabling this option (0).

ServerUUID

Description—Used internally by the Double-Take service to identify Double-Take connections and IP addresses used between servers

Values—Unique identifier generated by Double-Take

Default—Generated by Double-Take

Console Setting—None

Service restart required—Yes

Notes—If you are certain that the server is not being used by any jobs, you can delete the ServerUUID. For example, you may want to delete the ServerUUID so that you can create an image of a server after installing Double-Take. A deleted ServerUUID will be re-created the next time the Double-Take service is started. Keep in mind, if you delete the ServerUUID and the server is being used by any jobs, you will have problems with all aspects of Double-Take including mirroring, replication, and failover.

ServicePriority

Description—The priority level at which the Double-Take service runs.

Values—2 normal priority, 3 high priority

Default—2

Console Setting—None

Service restart required—Yes

Notes—The Double-Take service runs at normal priority by default. This option should not be modified, however, if the priority is raised to high (3), it can be done through Windows Task Manager.

ServicesToKeepRunning

Description—Services that will not be stopped on the target

Values—Semi-colon separated list of service names

Default—<null>

Console Setting—Set Options page, Target Services section, Services to leave running on the target server during protection

Service restart required—No

Notes—You can specify the service name using the service executable file name or the service display name. There is no need to use quotation marks, even if the names have spaces in them. Only separate the names by a semi-colon (;).

ServiceStopState

Description—Used internally by the Double-Take service. Do not modify this entry.

ShareFailover

Description—Specifies whether or not to failover shares

Values—0 Do not failover shares, 1 Failover shares

Default—1

Console Setting—Set Options page, Failover Options section, Failover shares

Service restart required—No

ShareUpdateInterval

Description—Specifies how often, in minutes, the share file will be sent to the target

Values—1 - 65535

Default—60

Console Setting—None

Service restart required—No

ShortFileNameScanIntervalMinutes

Description—Specifies how often, in minutes, the registry is scanned for short file names

Values—any valid integer

Default—240

Console Setting—None

Service restart required—No

ShutdownRebootTimeoutMinutes

Description—Specifies the amount of time, in minutes, to wait for the source to shutdown during failover or cutover

Values—any valid integer

Default—5

Console Setting—None

Service restart required—No

ShutdownTimeout

Description—The amount of time, in seconds, for the service to wait prior to completing the shutdown so that Double-Take can persist data on the target in an attempt to avoid a remirror when the target comes back online

Values—any valid number of seconds where 0 (zero) indicates waiting indefinitely and any other number indicates the number of seconds

Default—0

Console Setting—Edit Server Properties page, Setup section, Time allowed to complete shutdown operations

Service restart required—No

Notes—This setting only controls the service shutdown from the Double-Take clients. It does not control the service shutdown through a reboot or from the Service Control Manager.

SkipCompressionFileExt

Description—A period delimited list of file types that are not compressed, even if compression is enabled.

Values—any period delimited list of file types

Default—mp3.exe.wmv.wma.qt.mpg.mpeg.zip.jpg.jpeg.tiff.tar.rar.cab

Console Setting—None

Service restart required—No

SnapshotType

Description—Specifies the type of snapshot that Double-Take takes

Values—0 Create a client-accessible or non-client-accessible snapshot based on the job type , 1 Always create a client-accessible snapshot, 2 Always create a non-client-accessible snapshot

Default—0

Console Setting—None

Service restart required—No

SourceNewerMaxFileCount

Description—The number of files to compare during a source newer difference mirror

Values—1-1000

Default—16

Console Setting—None

Service restart required—No

SourcePendingAcks

Description—The number of operations received by the target queue in which the source is waiting for a response

Values—100 - 20,000

Default—2000

Console Setting—None

Service restart required—No

SourcePostFailbackScript

Description—Path on the source where the source post-failback script is located

Values—Any valid path

Default—<null>

Console Setting—None

Service restart required—No

SourcePostFailbackScriptArgs

Description—Arguments to be used with the source post-failback script

Values—Any valid argument

Default—<null>

Console Setting—None

Service restart required—No

SSMKeepTargetActivationCode

Description—Specifies if the activation code on the target is replaced or maintained after a full-serer failover or cutover. Do not modify this entry.

SSMShutdownServices

Description—Used by full server jobs to determine services to shutdown during failover or cutover. Do not modify this entry.

SSMShutdownSource

Description—Specifies if the source is shutdown when performing failover or cutover

Values—0 The source will not be shutdown, 1 The source will be shutdown

Default—1

Console Setting—Set Options page, Failover Options section, Shutdown the source server

Service restart required—Yes

Notes—This setting must be applied on the target server.

SSMStagingBase

Description—Specifies the folder to use for staging system state files for full server failover or cutover. Do not modify this entry.

SSMUseDiskSignature

Description—Used by full server jobs to determine how target disk signatures are used. Do not modify this entry.

StartupScript

Description—Used by full server jobs to control the post-failover script after reboot after failover. Do not modify this entry.

StatsDriverLogFlags

Description—Indicates which driver statistics are logged to the Double-Take log

Values—0 No driver statistics are logged, 1 State, 2 Operations, 4 Paging, 8 Timing

Default—0

Console Setting—None

Service restart required—Yes

Notes—Use the sum of various values to log multiple driver statistics. For example, a setting of 5 would log paging and state statistics. A setting of 7 would log paging, operations, and state statistics. A setting of 15 would log all driver statistics.

StatsFileName

Description—Default file for logging statistics

Values—any valid file name

Default—statistic.sts

Console Setting—Edit Server Properties page, Logging section, Filename (under Statistics)

Service restart required—No

StatsLoggingOn

Description—Specifies if Double-Take logs statistics at startup

Values—0 Stats logging does not start when Double-Take starts, 1 Stats logging starts when Double-Take starts

Default—0

Console Setting—Edit Server Properties page, Setup section, Setup Options, Log statistics automatically

Service restart required—No

StatsMaxFileSize

Description—Maximum size, in MB, for the statistic.sts file

Values—limited by available disk space

Default—10485760

Console Setting—Edit Server Properties page, Logging section, Maximum size (under Statistics)

Service restart required—No

StatsMaxObjects

Description—This entry is no longer used.

StatsPort

Description—Port used by pre-5.2 versions for DTStat to gather statistics

Values—1025 - 65535

Default—1106

Console Setting—None

Service restart required—Yes

StatsShmSize

Description—This entry is no longer used.

StatsWriteInterval

Description—Interval, in minutes, in which statistics are written to the statistic.sts file

Values—0 - 65535

Default—5

Console Setting—Edit Server Properties page, Logging section, Write interval

Service restart required—No

SystemMemoryLimit

Description—Set by the Double-Take service, each time it is started, to record the amount of available memory.

TargetPaused

Description—Internal setting that indicates if the target machine is paused. Do not modify this setting.

TargetPausedVirtual

Description—Internal setting that indicates which target machines are paused. Do not modify this setting.

TCPBufferSize

Description—Size of the TCP/IP buffer in bytes.

Values—4096-7500000

Default—375000

Console Setting—None

Service restart required—Yes

Notes—The default setting creates a TCP window that will accommodate most environments. In most environments, this value will not need to be adjusted. However, if your Double-Take network has a long end-to-end route and the throughput is not where you would expect it to be, then adjusting this parameter may have beneficial results. This value is the bandwidth delay product, which is calculated using the bandwidth of the network (in bits/second) times the round trip time (in seconds) between the two ends. Use the following recommended settings to improve Double-Take throughput performance.

- 100Mbit LAN—The setting should be around 37500.
- 1Gbit LAN—The setting should be around 375000.
- WAN—The setting should be around 130000.

While the calculations are fairly straight forward, the values that have been suggested are not exact because they depend on round trip time. Some improvements could be gained by adjusting these values either higher or lower. The value suited for your environment can best be determined through trial and error testing.

TempDir

Description—Temporary directory used when replicating Windows 200x encrypted files.

Values—Any valid path

Default—\Program Files\Vision Solutions\Double-Take\Temp

Console Setting—None

Service restart required—No

TGApplyMntPntSecurity

Description—Applies security settings to the volume of a mount point instead of applying them to the directory that the mount point is mounted to.

Values—0 Security will be applied to the directory, 1 Security will be applied to the volume

Default—0

Console Setting—None

Service restart required—Yes

Notes—This setting needs to be applied to the target server.

TGBlockOnConnect

Description—Blocks the target path for all connections, regardless of the source, so that the data cannot be modified

Values—0 Target paths are not blocked, 1 Target paths are blocked

Default—0

Console Setting—None

Service restart required—No

TGCloseDelay

Description—The length of time, in milliseconds, a file is held open on the target

Values—0 - 2000

Default—1000

Console Setting—None

Service restart required—No

Notes—If disk caching on the target is disabled either manually or by default (for example, by default on disks that host Active Directory database files), the target system may be slow during a mirror. If so, descreasing this setting to 100, 10, and 0 will result in incremental improvements, with 0 returning the system performance to normal.

TGDaysToKeepMovedFiles

Description—Specifies the length of time, in days, to keep moved files if TGMoveFilesOnDelete is enabled

Values—any valid integer

Default—0

Console Setting—Edit Server Properties page, Target section, Remove deleted files after this number of days

Service restart required—No

TGDisableAttributeReplication

Description—Specifies whether or not the attributes compression, ACL, and file mask are written to the target during mirroring and replication

Values—0 Enable attribute replication 1 Disable attribute replication

Default—0

Console Setting—None

Service restart required—No

TGExecutionRetryLimit

Description—The number of times an unfinished operation will be retried on the target before it is discarded. If this value is set to zero (0), an operation will never be discarded and will be retried on the target until it is applied.

Values—0 - 65536

Default—0

Console Setting—None

Service restart required—No

TGFileAlloc

Description—Indicates that Double-Take allocates an entire file on the first write of a mirror operation

Values—0 Disabled 1 Enabled

Default—1

Console Setting—None

Service restart required—No

Notes—To help eliminate file fragmentation on the target server, Double-Take should allocate the entire file first. With extremely large files, the file allocation may take a long time. Therefore, you may want to disable the file allocation. If you disable file allocation, you will have more fragmentation on the target disk.

TGMirrorCapacityHigh

Description—Maximum percentage of system memory that can contain mirror data before the target signals the source to pause the sending of mirror operations.

Values—2-75

Default—20

Console Setting—Edit Server Properties page, Target section, Pause mirroring at this level

Service restart required—No

TGMirrorCapacityLow

Description—Minimum percentage of system memory that can contain mirror data before the target signals the source to resume the sending of mirror operations.

Values—1-75

Default—15

Console Setting—Edit Server Properties page, Target section, Resume mirroring at this level

Service restart required—No

Notes—The maximum value forTGMirrorCapacityLow is either 75 or TGMirrorCapacityHigh, which ever is lower.

TGMoveFilesOnDelete

Description—Specifies whether files deleted on the source are actually moved to a different location on the target rather than being deleted on the target

Values—0 Files deleted on the source will be deleted on the target, 1 Files deleted on the source will be moved to a different location on the target

Default—0

Console Setting—Edit Server Properties page, Target section, Moved deleted files to this folder

Service restart required—No

Notes—If this option is enabled, the deleted files will be moved to the location specified in TGMoveFilesPath.

TGMoveFilesPath

Description—Specifies where deleted files on the source are being moved to on the target

Values—any valid path

Default—<null>

Console Setting—Edit Server Properties page, Target section, Moved deleted files to this folder

Service restart required—No

TGMoveFilesSingleDirectory

Description—Specifies if deleted files that will be moved on the target (see **TGMoveFilesOnDelete**) will be moved to a single directory structure

Values—0 Use the same directory structure on the target as the source to store deleted files, 1 Use a single directory structure on the target to store deleted files

Default—0

Console Setting—None

Service restart required—No

TGRetryLocked

Description—Minimum number of seconds to wait before retrying a failed operation on a target

Values—0-65536

Default—3

Console Setting—Edit Server Properties page, Target section, Retry delay for incomplete operations

Service restart required—No

TGThreadCount

Description—This setting is no longer used

TGUnfinishedOpEvent

Description—Specifies whether or not unfinished operations on the target are logged to the Event Viewer

Values—0 Unfinished operation messages are not logged, 1 Unfinished operation messages are logged

Default—1

Console Setting—None

Service restart required—No

TGWriteCache

Description—Specifies whether or not Double-Take uses the intermediate cache

Values—0 Bypass the intermediate cache and write directly to disk, 1 Do not bypass the intermediate cache

Default—0 for full server to ESX appliance jobs, 1 for all other job types

Console Setting—None

Service restart required—No

TGWriteFailureBeforeNotification

Description—Specifies the number of times an operation will be retried on the target before a notification is sent to update the target status

Values—0-1024

Default—10

Console Setting—None

Service restart required—Yes

Notes—If you change the setting to 0, the notification will be disabled. Changing this option will only affect how the target status is displayed. To solve the underlying issue of why the operations are failing will require investigation into the Double-Take log files.

UpdateInterval

Description—This setting is no longer used

UpgradeCode

Description—Used by the Double-Take installation program to maintain the installation settings for an upgrade. Do not modify this entry.

UseChangeJournal

Description—Specifies if the Windows NTFS change journal is used to track file changes. If the source is rebooted, only the files identified in the change journal will be remirrored to the target. This setting helps improve mirror times.

Values—0 Do not track file changes, 1 Track file changes and remirror only changed files on source reboot, 2 Track file changes and remirror only changed files on source reboot and auto-reconnect

Default—1

Console Setting—Edit Server Properties page, Setup section, Mirror only changed files when source reboots

Service restart required—Yes

Notes—The corresponding Console Setting only allows off (0) and on (1) settings. Therefore, if you set UseChangeJournal to 2, the corresponding Console Setting will be disabled. The Console Setting will be reenabled when UseChangeJournal is set to 0 or 1.

UseEventLog

Description—Specifies whether or not messages are logged to the Windows Event Viewer

Values—0 Do not log messages to the Event Viewer, 1 Log messages to the Event Viewer

Default—1

Console Setting—None

Service restart required—No

UseLegacyDrivers

Description—This setting is no longer used

UserIntervention

Description—Specifies whether or not user intervention is required to initiate a failover or cutover

Values—0 User intervention is not required, 1 User intervention is required

Default—1

Console Setting—Set Options page, Failover Options section, Wait for user to initiate failover

Service restart required—No

UseScheduledPause

Description—Used by Double-Take for internal schedule processing. Do not modify this setting.

UseShareFile

Description—Specifies whether to create and use a share file or to use the shares that are currently stored in the target memory

Values—0 Use the shares that are currently stored in the target memory, 1 Create and use a file containing the share information

Default—1

Console Setting—None

Service restart required—No

VerifyLogAppend

Description—Specifies whether the DTVerify.log file will be appended to or overwritten

Values—0 Overwrite, 1 Append

Default—1

Console Setting—Edit Server Properties page, Logging section, Append

Service restart required—No

VerifyLogLimit

Description—Maximum size of the DTVerify.log file in bytes

Values—limited by available hard drive space, up to 4 GB

Default—1048576

Console Setting—Edit Server Properties page, Logging section, Maximum size (under Verification)

Service restart required—No

VerifyLogName

Description—Name of the verification log file

Values—any valid file name

Default—DTVerify.log

Console Setting—Edit Server Properties page, Logging section, File name (under Verification)

Service restart required—No

VerifyRetryInterval

Description—The time, in minutes, between when one verification fails and a retry is scheduled to begin.

Values—any valid number

Default—3

Console Setting—None

Service restart required—No

VerifyRetryLimit

Description—The number of time a verification will be retried.

Values—any valid number

Default—5

Console Setting—None

Service restart required—No

VersionInfo

Description—The version of Double-Take that was installed. Do not modify this entry.

WarningPings

Description—This entry is no longer used.

WatchDogFailureProcessDump

Description—Creates a troubleshooting dump file if the Double-Take driver stops running

Values—0 Do not create a dump file, 1 Create a dump file

Default—0

Console Setting—None

Service restart required—No

WatchDogFailureScript

Description—Specifies the script to run if the Double-Take driver stops running

Values—Any valid path and script file

Default—<null>

Console Setting—None

Service restart required—No

Viewing server events

Highlight a server on the **Manage 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 *Event messages* on page 718 for a complete list of the service and driver event messages.

Viewing server logs

You can view the Double-Take and Double-Take Management Service log files through the Double-Take Console using either of these two methods.

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

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

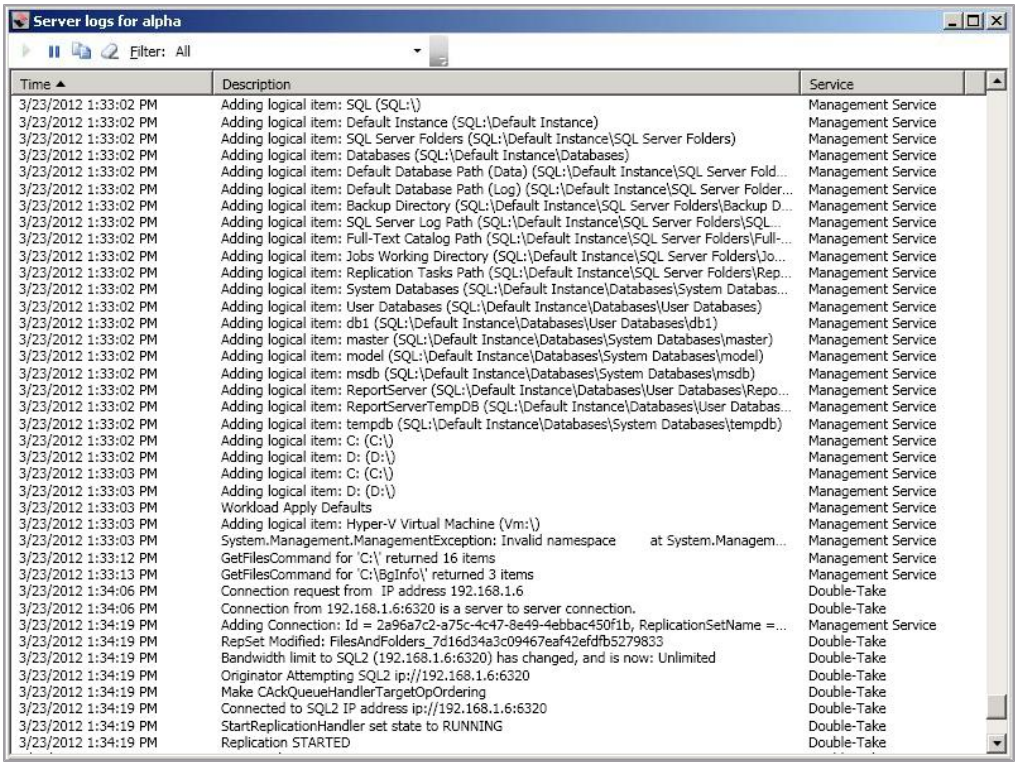

The following table identifies the controls and the table columns in the **Server logs** window.

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.

See *Log files* on page 677 for more information on the log files.

Managing VMware servers

To manage your VMware servers, select **Go**, **Manage VMware Servers**. The **Manage VMware Server** page allows you to view, add, remove, or edit credentials for your VMware servers available in the console.

VMware Server

The name of the VMware server

Full Name

The full name of the VMware server

User Name

The user account being used to access the VMware server

Add VMware Server 

Add a new WMare server. When prompted, specify the VMware server and a user account.

Remove Server 

Remove the VMware server from the console.

Provide Credentials 

Edit credentials for the selected VMware server. When prompted, specify a user account to access the VMware server.

Managing snapshots

A snapshot is an image of the source replica data on the target taken at a single point in time. Snapshots allow you to view files and folders as they existed at points of time in the past, so you can, for example, recover from cases where corrupted source data was replicated to the target. For some Double-Take job types, when failover is triggered, you can use the live target data at the time of failover or you can failover to a snapshot of the target data.

1. From the **Manage 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.
   - **User Request**—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.

Chapter 6 Selecting a protection type

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

- **Data protection**—If you want to protect specific data on your source, your job type will be a files and folders job. Review the following data decision tree to see the various configurations available for a files and folder job.

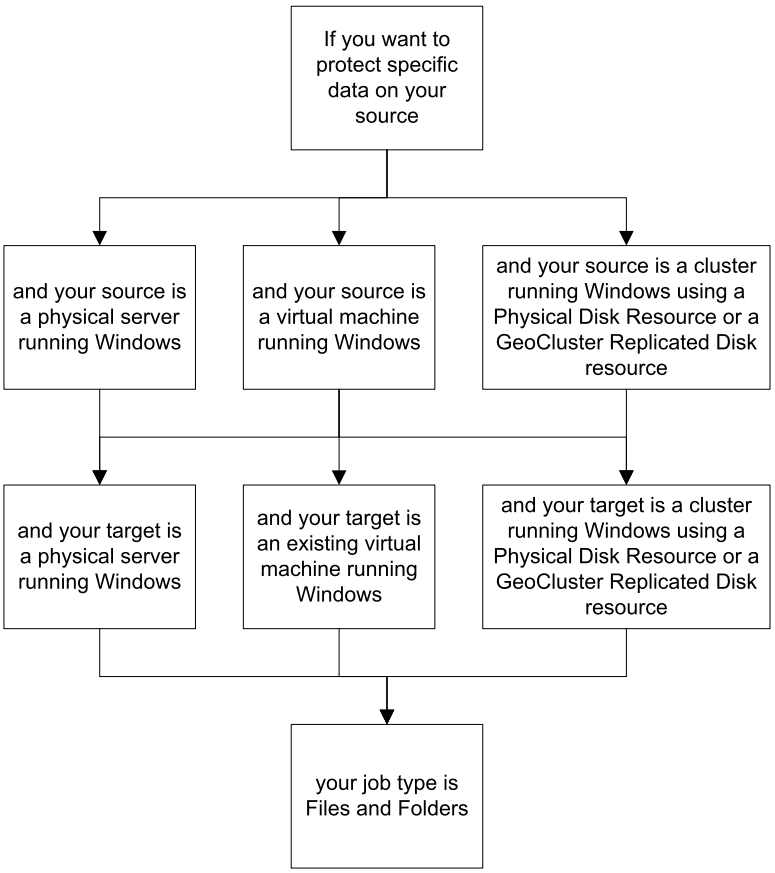

If your source is a domain controller, you should use one of the full server protection methods to protect the entire server because of complexities with SPNs.

- **Application protection**—If you want to protect Microsoft Exchange or Microsoft SQL, your job type will be an Exchange Server or SQL Server job. Review the following application decision tree to see the various configurations available for these job types.

If your source is a domain controller, you should use one of the full server protection methods to protect the entire server because of complexities with authentication and as a best practice.

- **Full server protection**—Protecting an entire server is the most flexible, but confusing, scenario. Unlike the data and application decision trees, the full server decision trees leads you to many different possible Double-Take job types. Because the full server decision tree is so large and complex, it is broken down into smaller sections on the following pages. Study the full server decision tree, or the following smaller sections, to determine the best job type for your environment and protection needs.

Physical source server branches of the full server decision tree

Windows at the virtual guest level of an ESX host branches of the full server decision tree

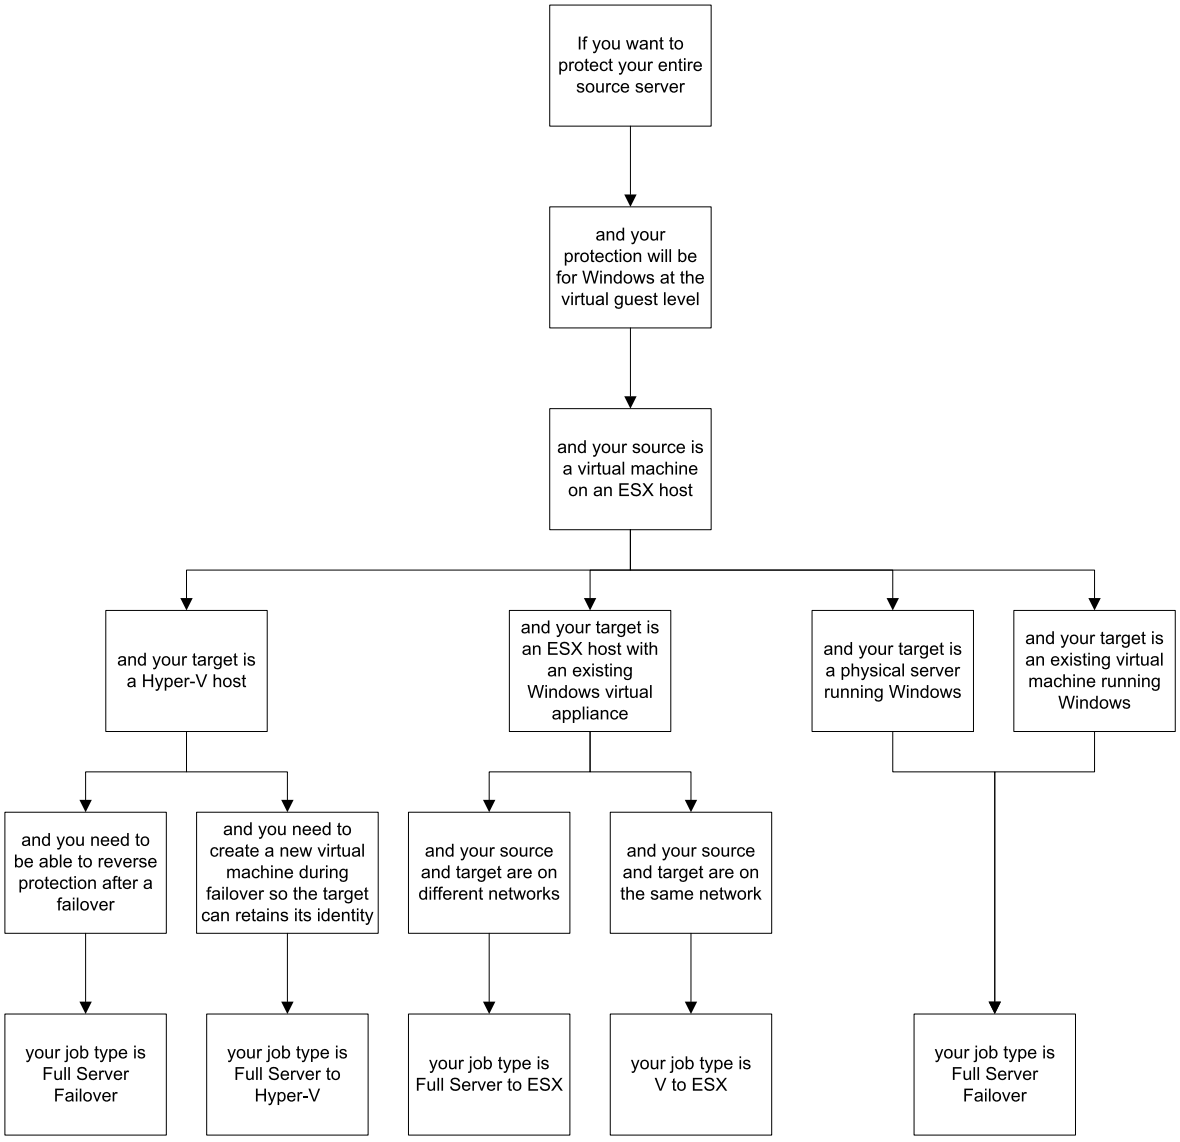

Windows at the virtual guest level of a Hyper-V host branches of the full server decision tree

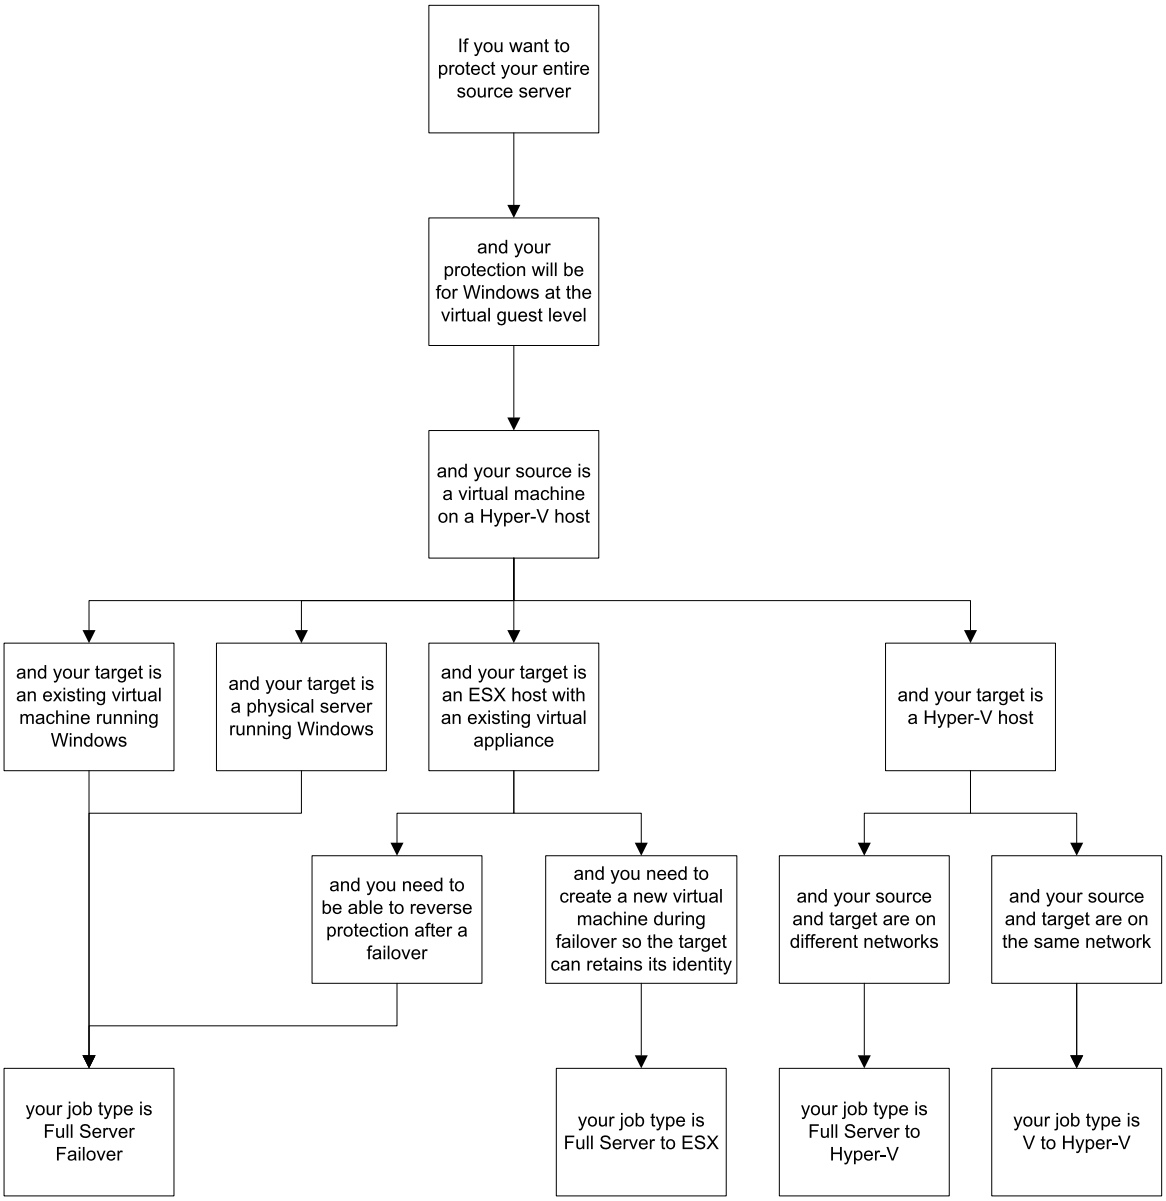

Windows at the virtual guest level of non-ESX or non-Hyper-V host branches of the full server decision tree

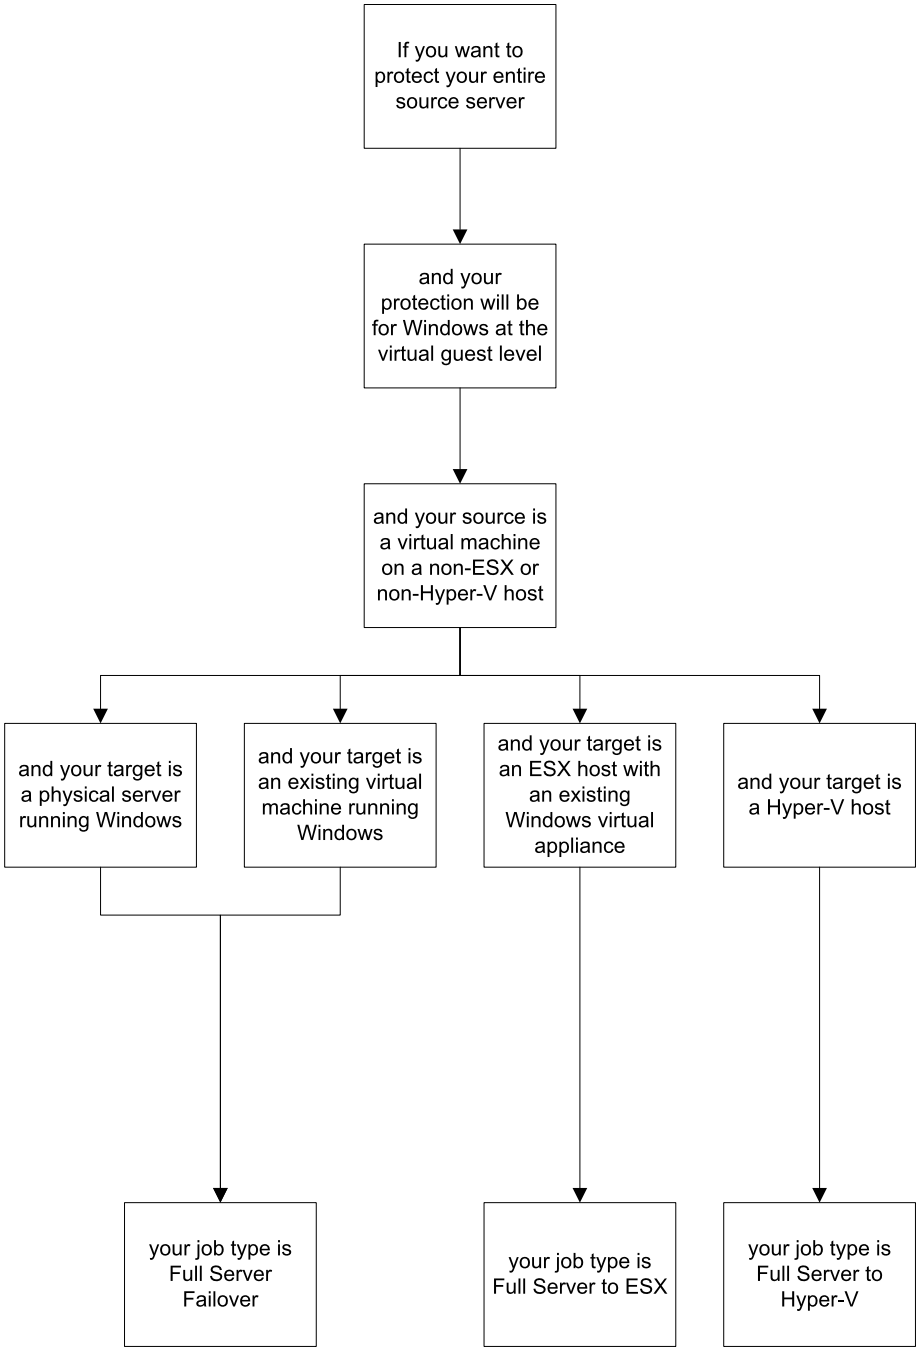

Host level branches of the full server decision tree

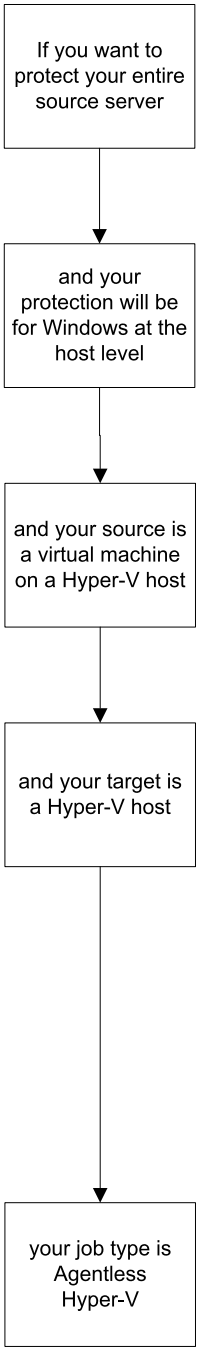

Chapter 7 Files and folders protection

Create a files and folders job when you want to protect data or file shares. You can also use it to protect applications, such as Oracle or MySQL, however you will need to use your own customized failover and failback scripts to start and stop services during failover and failback. This job type does not protect a server's system state.

- See *Files and folders requirements* on page 174—Files and folders protection includes specific requirements for this type of protection.
- See *Creating a files and folders job* on page 176—This section includes step-by-step instructions for creating a files and folders job.
- See *Managing and controlling files and folders jobs* on page 207—You can view status information about your files and folders jobs and learn how to control these jobs.
- See *Failing over files and folders jobs* on page 225—Use this section when a failover condition has been met or if you want to failover manually.
- See *Failback and restoration for files and folders jobs* on page 226—Use this section to determine if you want to failback and then restore or if you want to restore then failback.

If your source is a domain controller, you should use one of the full server protection methods to protect the entire server because of complexities with SPNs. See *Selecting a protection type* on page 165.

Files and folders requirements

In addition to the *Core Double-Take requirements* on page 23, use these requirements for full server protection.

- **Microsoft Server Core 2008 R2, 2012, or 2012 R2**—These operating systems are supported for mirroring, replication, and failover. However, DNS updates are not supported for Server Core servers.
- **Snapshots**—Support for Double-Take snapshots for files and folders jobs is limited. You can take snapshots, however you cannot failover to a snapshot. The snapshots can be accessed and reverted on the target manually using VSS tools and utilities. Additionally, snapshots are not supported if your source and/or target is a cluster. See *Core Double-Take requirements* on page 23 for the specific snapshot requirements.
- **NAT support**—Files and folders jobs (non-cluster) can support NAT environments in an IP-forwarding configuration with one to one port mappings. Port-forwarding is not supported. Additionally, only IPv4 is supported for NAT environments. Make sure you have added your servers to the Double-Take Console using the correct IP address. Review the *NAT configuration* table on page 72 in the *Adding servers* section before you start the job creation process.
- **Supported configurations**—The following table identifies the supported configurations for a files and folders job.

Configuration		Supported	Not Supported
Source to target configuration[1]	One to one, active/standby	X	
	One to one, active/active	X	
	Many to one	X	
	One to many	X	
	Chained	X	
	Single server	X	
Server configuration	Standalone to standalone	X	
	Standalone to cluster [2]	X	
	Cluster to standalone	X	
	Cluster to cluster	X	
	Cluster Shared Volumes (CSV) guest level	X	
	Cluster Shared Volumes (CSV) host level		X

Configuration		Supported	Not Supported
Upgrade configuration[3]	Upgrade 5.3 Replication Console connection to 7.0 Double-Take Console files and folders job		X
	Upgrade 5.3 Double-Take Source Connection resource connection to 7.0 Double-Take Console files and folders cluster job		X
	Upgrade 5.3 files and folders job to 7.0 files and folders job	X	
	Upgrade 6.0 files and folders job to 7.0 files and folders job	X	
Version 7.0 console[4]	Version 7.0 console can create job for 5.3 source and 5.3 target	X	
	Version 7.0 console can create job for 6.0 source and 6.0 target	X	
	Version 7.0 console can create job for 7.0 source and 7.0 target	X	

1. See *Supported configurations* on page 16 for details on each of the source to target configurations.
2. Standalone to cluster configurations do not support failover.
3. When using a supported upgrade configuration, you can perform a rolling upgrade where you update the target server first. After the upgrade is complete, the source will automatically reconnect to the target. Upgrade the source when convenient. For an upgrade configuration that is not supported, you will have to delete the existing connection before the upgrade and create a new job after the upgrade.
4. Newer job options available in the version 7.0 console will not be functional when creating jobs for servers running version 5.3 or 6.0.

Creating a files and folders job

Use these instructions to create a files and folders job.

With a file and folders job (non-cluster), your servers can be in a NAT environment. However, you must make sure you have added your servers to the Double-Take Console using the correct IP address. Review the *NAT configuration* table on page 72 in the *Adding servers* section before you start the job creation process.

Note the following information concerning cluster scenarios.

- If your source is a cluster, it must have a file server role or resource configured in the group that contains the disk you are wanting to protect
- For cluster to standalone jobs, the source cluster name and the standalone target must be added to the **Manage Servers** page. Double-Take will query the source cluster name and get a list of cluster groups that contain a file share resource and then present those groups for protection. Double-Take will failover the name, IP addresses, SPNs, and shares for the clustered file server to the standalone target.
- For cluster to cluster jobs, the source and target cluster names must be added to the **Manage Servers** page. Double-Take will query the source cluster name and get a list of cluster groups that contain a file share resource and then present those groups for protection. Double-Take will not failover the name, IP addresses, SPNs, or shares. Failover will be DNS failover. Groups will be brought online on the target, therefore they need to be pre-staged using step 1 or 2 below, depending on your operating system.
- For standalone to cluster jobs, the standalone source and the target cluster name must be added to the **Manage Servers** page. Double-Take will query the target cluster name and determine which cluster group contains the disk resource where data will be sent. Failover is not supported for standalone to cluster files and folders jobs. Jobs in this configuration can only be used for data protection or migration, not high availability.

1. If you are using a cluster to cluster configuration and you are using Windows 2003, follow these instructions to configure the target cluster before you create your Double-Take files and folders job.
    a. Take the source file server Network Name offline.
    b. Create the same Network Name on the target cluster in the file server group.
    c. Create identical file shares in the target file server group and make them dependent on the Network Name that was just created.
    d. Take the target Network Name resource offline.
    e. Bring the source Network Name resource online.
    f. If it is not already, DNS registration must be enabled for a source cluster name for DNS failover to function properly. You should not allow the source cluster name to come online if your source fails over.
2. If you are using a cluster to cluster configuration and you are using Windows 2008 or 2012, follow these instructions to configure the target cluster before you create your Double-Take files and folders job.

a. You may want to lower the TTL (time to live) value of your DNS records to avoid caching during this configuration.
b. Give the target cluster account full control over the source file server account in Active Directory and over the source file server DNS record.

If control of the DNS record is not added, the Client Access Point will fail name resolution and you will not be able to create the target shares unless the record is manually updated.

c. Take the source file server Client Access Point offline, including IP address, and then flush DNS on the target.
d. Create the Client Access Point with the same name on the target and provide a unique IP address.
e. While the target Client Access Point is online, create identical shares to the source file server.
f. Take the target Client Access Point and IP address offline. Take the target file share resources offline.
g. Bring the source Client Access Point online. DNS registration will change the IP address back to the source IP address.
h. If desired, reset your TTL value back to its original value.
i. If it is not already, DNS registration must be enabled for a source cluster name for DNS failover to function properly. You should not allow the source cluster name to come online if your source fails over.

3. Click **Get Started** from the toolbar.
4. Select **Double-Take Availability** and click **Next**.
5. Select **Protect files and folders, an application, or an entire Windows server** and click **Next**.
6. Choose your source server. This is the server that contains the files and folders that you want to protect. If your source is a cluster, select the source cluster name, not the file server or node name.

- **Current Servers**—This list contains the servers currently available in your console session. Servers that are not licensed for the workflow you have selected will be filtered out of the list. Select your source server from the list.
- **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 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.

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. Ideally, you should use a local account rather than a domain account because the domain account will fail to authenticate while failed over if the NetBIOS name and/or SPNs are failed over. If you want Double-Take to update DNS during failover, the account must be a member of the Domain Admins group. If your security policies do not allow use of this group, see *DNS* on page 816 and use the instructions under the Double-Take DFO utility to use a non-Domain Admins account.

To best prepare for a potential failover, restore, and failback scenario, you may want to use a source server that has been inserted into the console by a private IP address. In this scenario, your source must have at least two IP addresses, one for public communication and one for private. If you insert your source in the console using the private IP address, then that private IP address can more easily be used after a failure to restore the data that changed during failover from the target back to the source. If your source server has

already been inserted into the console by name, you can remove it and reinsert it by private IP address on the **Manage Servers** page.

7. Click **Next** to continue.
8. Choose the type of workload that you want to protect. Under **Server Workloads**, in the **Workload types** pane, select **Files and Folders**. In the **Workload items** pane, you will see the volumes and shares (if any) for your source, or if you are protecting a cluster, you will see your cluster groups. Select the volumes, shares, or cluster groups that you want to protect. You can select your files and folders in more detail in the **Replication Rules** section.

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

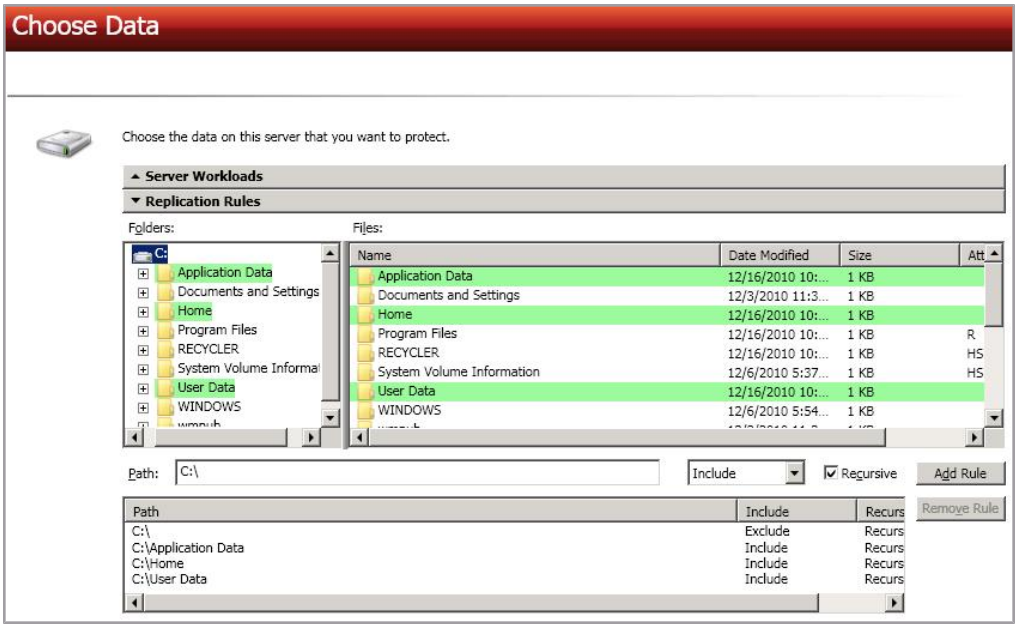

Volumes and folders with a green highlight are included completely. Volumes and folders highlighted in light yellow are included partially, with individual files or folders included. If there is no highlight, no part of the volume or folder is included. To modify the items selected, highlight a volume, 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.

---

10. Click **Next** to continue.
11. Choose your target server. This is the server that will store the replica data from the source.If your target is a cluster, select the target cluster name, not the file server or node name.

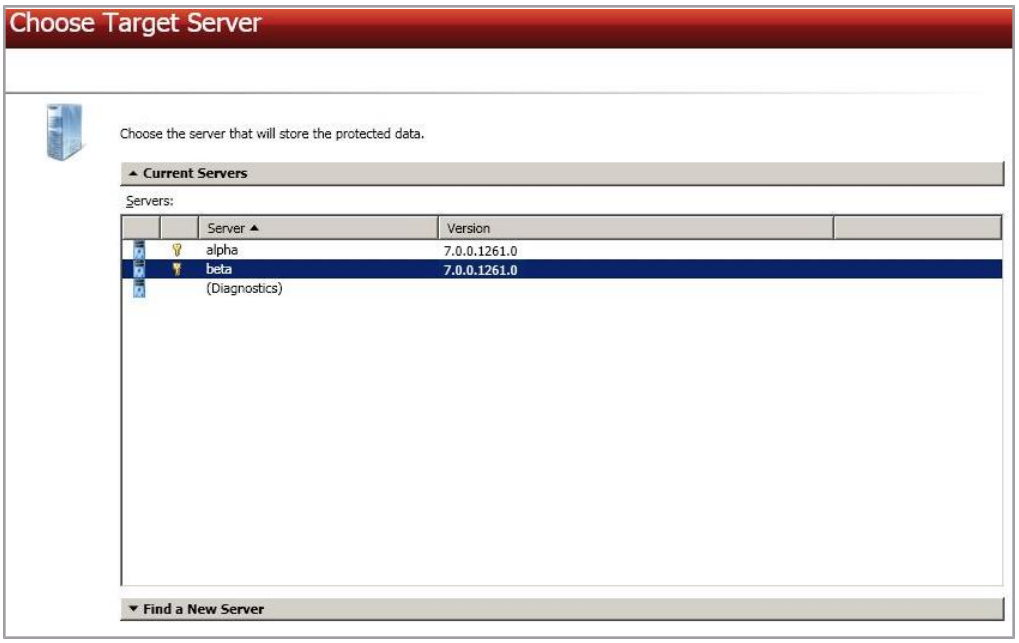

- **Current Servers**—This list contains the servers currently available in your console session. Servers that are not licensed for the workflow you have selected and those not applicable to the workload type you have selected will be filtered out of the list. Select your target server from the list.
- **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. Ideally, you should use a local account rather than a domain account because the domain account will fail to authenticate while failed over if the NetBIOS name and/or SPNs are failed over. If you want Double-Take to update DNS during failover, the account must be a member of the Domain Admins group. If your security policies do not allow use of this group, see *DNS* on page 816 and use the instructions under the Double-Take DFO utility to use a non-Domain Admins account.

12. Click **Next** to continue.
13. You have many options available for your files and folders job. However, not all of the options will be applicable to clustered environments or to single server configurations. 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 206.

    - *General* on page 182
    - *Failover Monitor* on page 183
    - *Failover Options* on page 186
    - *Failover Identity* on page 189
    - *Mirror, Verify & Orphaned Files* on page 192
    - *Target Route* on page 196
    - *Network Route* on page 197
    - *Path Mapping* on page 198
    - *Failover Services* on page 200
    - *Snapshots* on page 201
    - *Compression* on page 202
    - *Bandwidth* on page 203
    - *Scripts* on page 205

General

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

Failover Monitor

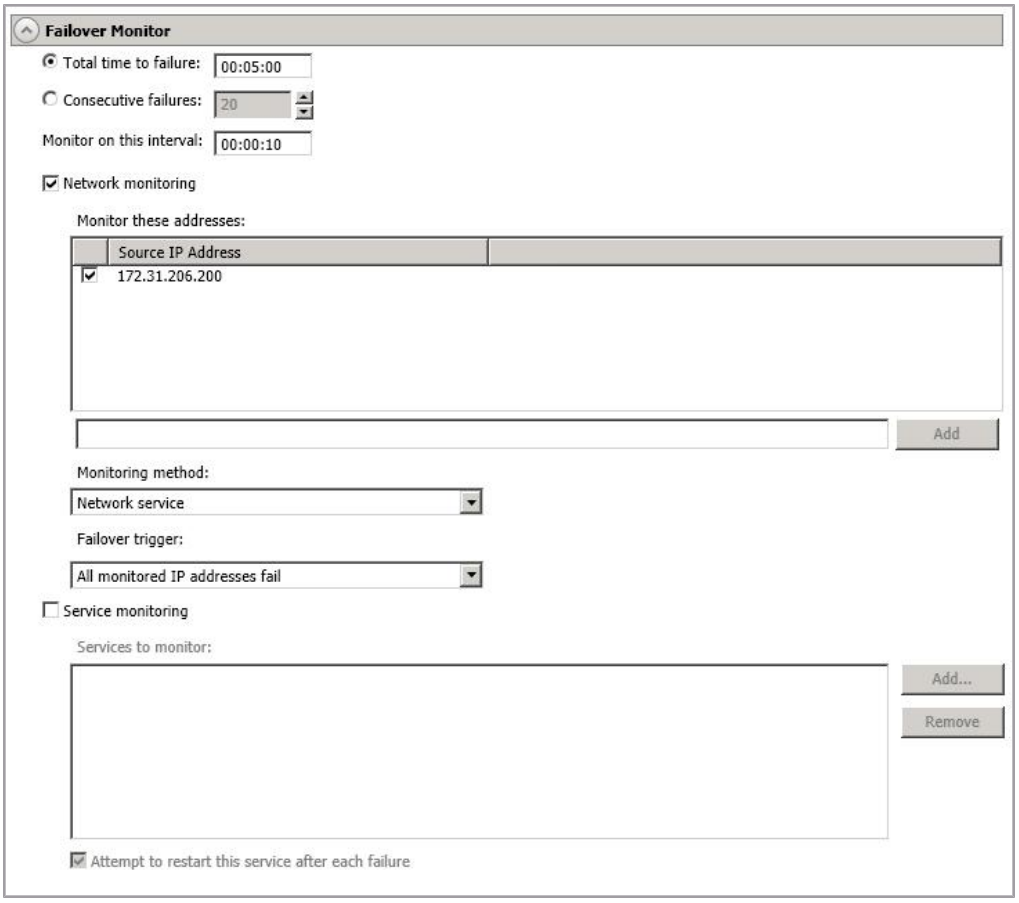

- **Total time to failure**—Specify, in hours:minutes:seconds, how long the target will keep trying to contact the source before the source is considered failed. This time is precise. If the total time has expired without a successful response from the source, this will be considered a failure.

  Consider a shorter amount of time for servers, such as a web server or order processing database, which must remain available and responsive at all times. Shorter times should be used where redundant interfaces and high-speed, reliable network links are available to prevent the false detection of failure. If the hardware does not support reliable communications, shorter times can lead to premature failover. Consider a longer amount of time for machines on slower networks or on a server that is not transaction critical. For example, failover would not be necessary in the case of a server restart.

- **Consecutive failures**—Specify how many attempts the target will make to contact the source before the source is considered failed. For example, if you have this option set to 20, and your source fails to respond to the target 20 times in a row , this will be considered a failure.
- **Monitor on this interval**—Specify, in hours:minutes:seconds, how long to wait between attempts to contact the source to confirm it is online. This means that after a response (success or failure) is received from the source, Double-Take will wait the specified interval

time before contacting the source again. If you set the interval to 00:00:00, then a new check will be initiated immediately after the response is received.

If you choose **Total time to failure**, do not specify a longer interval than failure time or your server will be considered failed during the interval period.

If you choose **Consecutive failures**, your failure time is calculated by the length of time it takes your source to respond plus the interval time between each response, times the number of consecutive failures that can be allowed. That would be (response time + interval) * failure number. Keep in mind that timeouts from a failed check are included in the response time, so your failure time will not be precise.

- **Network monitoring**—With this option, the target will monitor the source using a network ping. Disable this option if you are using a standalone to cluster configuration, because failover is not supported.
  - **Monitor these addresses**—Select each **Source IP Address** that you want the target to monitor. If you want to monitor additional addresses, enter the address and click **Add**. In a NAT environment, you can add any additional public IP addresses on the source that may not be listed. Additionally, you should not monitor any private IP addresses on the source because the target cannot reach the source's private address in a NAT environment, thus causing an immediate failure.

    

    If you are protecting a cluster, you are limited to the IP addresses in the cluster group that you are protecting.

  - **Monitoring method**—This option determines the type of network ping used for failover monitoring.
    - **Network service**—Source availability will be tested by an ICMP ping to confirm the route is active.
    - **Replication service**—Source availability will be tested by a UDP ping to confirm the Double-Take service is active.
    - **Network and replication services**—Source availability will be tested by both an ICMP ping to confirm the route is active and a UDP ping to confirm the Double-Take service is active. Both pings must fail in order to trigger a failover.
  - **Failover trigger**—If you are monitoring multiple IP addresses, specify when you want a failover condition to be triggered.
    - **One monitored IP address fails**—A failover condition will be triggered when any one of the monitored IP addresses fails. If each IP address is on a different subnet, you may want to trigger failover after one fails.
    - **All monitored IP addresses fail**—A failover condition will be triggered when all monitored IP addresses fail. If there are multiple, redundant paths to a server, losing one probably means an isolated network problem and you should wait for all IP addresses to fail.
- **Service monitoring**—This option is only available in standalone environments. It is not available in cluster environments.With this option, the target will monitor specific services on the source by confirming that they are running. Multiple services in the list will be checked in parallel. A failover condition is met when one of the monitored services fails the

check. Click **Add** and select the service that you want to monitor. Repeat this step for additional services that you want to monitor. If you want to remove a service from the **Services to monitor** list, highlight it and click **Remove**.

- **Attempt to restart this service after each failure**—When this option is enabled, if a service fails the monitor check, Double-Take will attempt to restart it. During this restart period, Double-Take will not check the other services, to avoid any false failures while the one service is attempting to be restarted. If the service cannot be restarted, Double-Take will consider this a failure.

Failover Options

If you are using a standalone to cluster configuration, you can skip this section because failover is not supported.

If you want to disable failover monitoring completely, you must disable **Failover shares**, **Failover host name**, and **Failback host name**, and do not specify anything for the failover and failback scripts. Additionally, you must select **Retain target network configuration** in the **Failover Identity** section. If you select any of these options or choose **Apply source network configuration to the target** under **Failover Identity**, a failover monitor will be created.

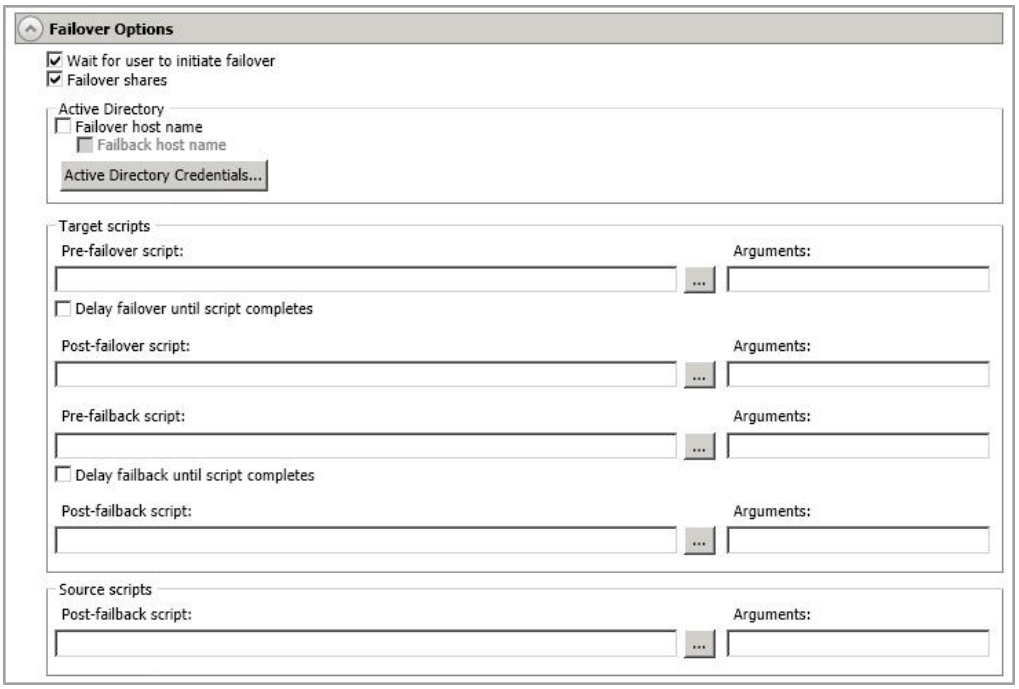

- **Wait for user to initiate failover**—By default, the failover process will wait for you to initiate it, allowing you to control when failover occurs. When a failure occurs, the job will wait in **Failover Condition Met** for you to manually initiate the failover process. Disable this option only if you want failover to occur immediately when a failure occurs.
- **Failover shares**—By default, shares will be failed over to the target. If you do not want to failover shares, disable this option.

  Automatic share failover only occurs for standard Windows file system shares. Other shares must be configured for failover through the failover scripts or created manually on the target. See *Macintosh shares* on page 826 or *NFS Shares* on page 827 for more information.

  If you are failing over Windows shares but your source and target do not have the same drive letters, you must use the **All to One** selection under **Path Mapping**

when establishing your job. Otherwise, the shares will not be created on the target during failover.

If your target is a standalone server, Windows share information is automatically updated on the target once an hour. Since shares are not failed over, if your target is a cluster, you will need to manually update shares on a cluster target. For Windows 2003, you can create the file share resources while the source is running but keep the resource offline. For Windows 2008 and 2012, you will need to repeat the initial cluster configuration steps, although you can skip giving the target cluster account full control and creating the target Client Access Point, because those steps were already completed during the initial cluster configuration.

- **Failover host name**—If desired, you can failover the source server's host name. This will automatically remove the host SPN (Service Principle Name) from Active Directory on the source and add it to Active Directory on the target. If you are using Active Directory, enable this option or you may experience problems with failover.
- **Failback host name**—This option returns the host SPN on the source and target back to their original settings on failback. If you are using Active Directory, enable this option or you may experience problems with failback.
- **Active Directory Credentials**—If you are failing over and/or failing back the host name, you need to specify a user that has update privileges within Active Directory. Click **Active Directory Credentials** and identify a user and the associated password that has privileges to create and delete SPNs. The username must be in the format fully_qualified_ domain\user,, and the account password cannot be blank.
- **Scripts**—You can customize failover and failback by running scripts on the source and target. 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 99. Examples of functions specified in scripts include stopping services on the target before failover because they may not be necessary while the target is standing in for the source, stopping services on the target that need to be restarted with the source's machine name and/or IP address, starting services or loading applications that are in an idle, standby mode waiting for failover to occur, notifying the administrator before and after failover or failback occurs, stopping services on the target after failback because they are no longer needed, stopping services on the target that need to be restarted with the target machine's original name and/or IP address, and so on. There are four types of failover and failback scripts.
  - **Pre-failover script**—This script runs on the target at the beginning of the failover process. Specify the full path and name of the script file.
  - **Post-failover script**—This script runs on the target at the end of the failover process. Specify the full path and name of the script file.
  - **Pre-failback script** —This script runs on the target at the beginning of the failback process. Specify the full path and name of the script file.
  - **Post-failback script**—This script runs on the target or source at the end of the failback process. Specify the full path and name of the script file.
  - **Arguments**—Specify a comma-separated list of valid arguments required to execute the script.

- **Delay until script completes**—Enable this option if you want to delay the failover or failback process until the associated script has completed. If you select this option, make sure your script handles errors, otherwise the failover or failback process may never complete if the process is waiting on a script that cannot complete.

Scripts will run but will not be displayed on the screen if the Double-Take service is not set to interact with the desktop. Enable this option through the Windows Services applet.

Failover Identity

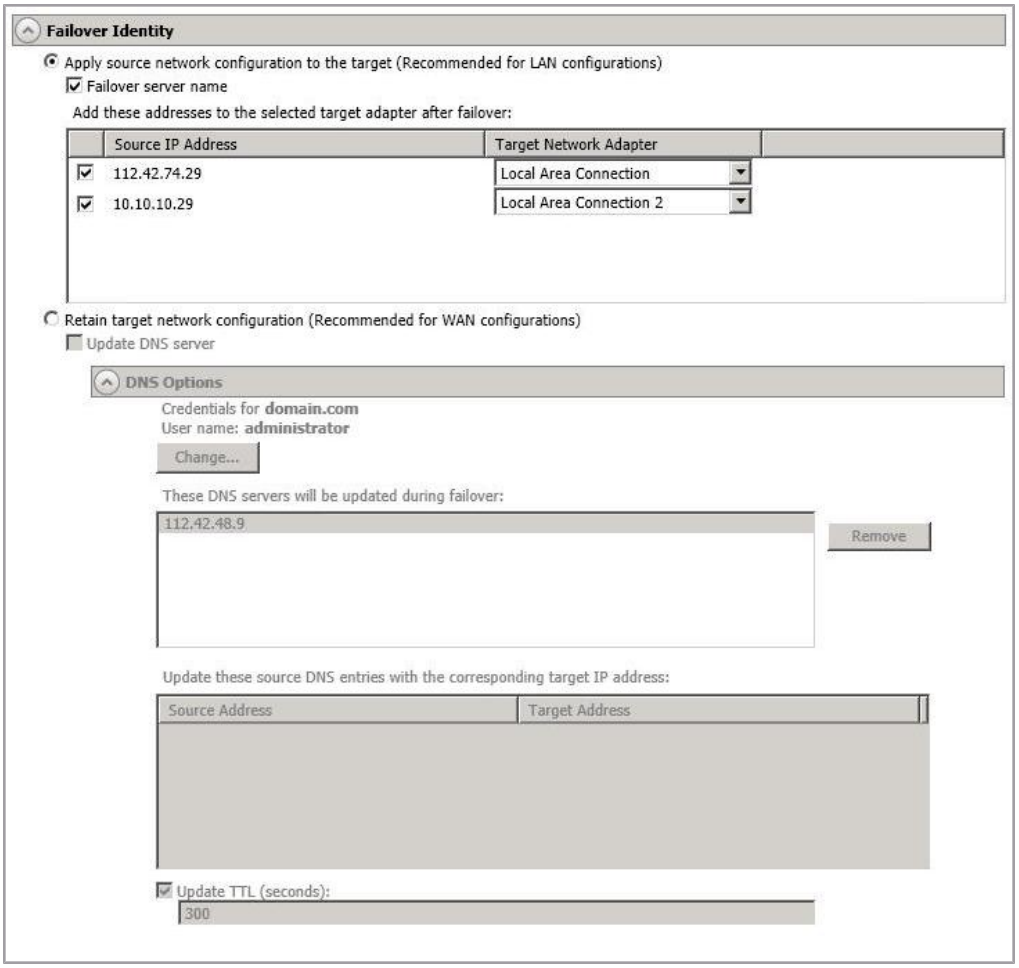

If you want to disable the ability to failover, you must select **Retain target network configuration**. Additionally, in the **Failover Options** section, you must disable **Failover shares**, **Failover host name**, and **Failback host name**, and do not specify anything for the failover and failback scripts. If you select any of these options or choose **Apply source network configuration to the target**, the ability to failover will not be disabled.

- **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 NAT router.

  If you are applying the source network configuration to the target in a WAN environement, do not failover your IP addresses 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.

- **Failover server name**—Select this option to failover the server name to the target. Double-Take checks the hosts file and uses the first name there. If there is no hosts file, Double-Take will use the first name in DNS. (Keep in mind, the first name in DNS may not always be the same each time the DNS server is rebooted.) Lastly, if there is no DNS server, Double-Take will use the failover monitor name created by the Double-Take Console. You will need to failover the server name for file shares to function properly after failover. If you do not want to failover the server name, disable this option.
- **Add these addresses to the selected target adapter after failover**—Select which IP addresses you want to failover and select the **Target Network Adapter** that will assume the IP address during failover.

If you have inserted your source server into the console using a private IP address, do not select the private IP address for failover.

If you configured failover to be triggered when all monitored IP addresses fail and are failing over more IP addresses than you are monitoring, you may have IP address conflicts after failover. For example, if you monitor two out of three addresses, and those two addresses fail but the third one does not, and you failover all three IP addresses, the third address that did not fail may exist on both the source and the target, depending on the cause of the failure. Therefore, when a source is failing over more IP addresses than are being monitored, there is a risk of an IP address conflict.

If you are protecting a cluster, you are limited to the IP addresses in the cluster group that you are protecting.

- **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.

If you choose to retain the target network configuration, but have also selected to failover shares under the **Failover Options** sections, the source NetBIOS name will automatically be failed over so that the shares can be accessed after failover.

- **Update DNS server**—Specify if you want Double-Take to update your DNS server on failover. If DNS updates are made, the DNS records will be locked during failover. Be sure and review the *Core Double-Take requirements* on page 23 for the requirements for updating DNS.

  

  DNS updates are not available for Server Core servers or NAT configurations.

  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. For clusters, be sure and select the clustered IP address.
  - **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.

If you select **Retain your target network configuration** but do not enable **Update DNS server**, you will need to specify failover scripts that update your DNS server during failover, or you can update the DNS server manually after failover. This would also apply to non--Microsoft Active Directory Integrated DNS servers. You will want to keep your target network configuration but do not update DNS. In this case, you will need to specify failover scripts that update your DNS server during failover, or you can update the DNS server manually after failover.

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

Mirror, Verify & Orphaned Files

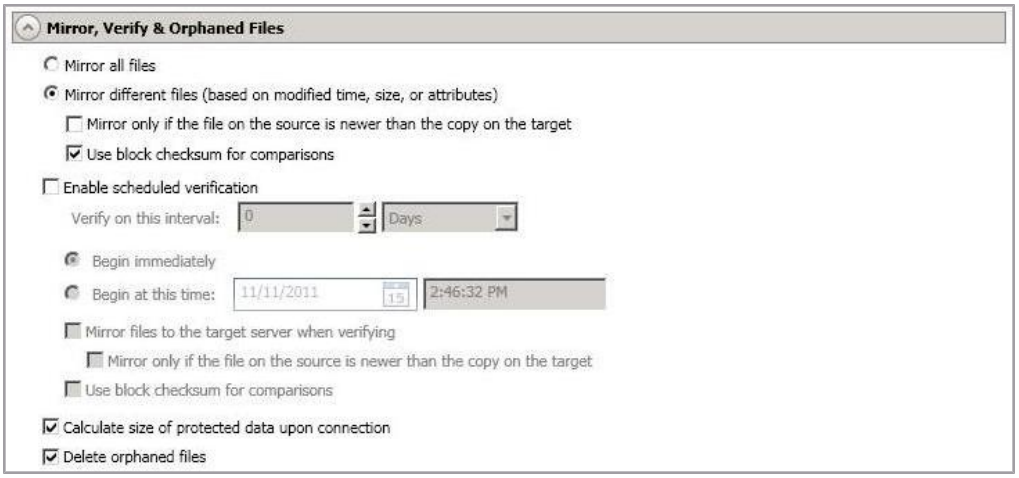

- **Mirror all files**—All protected files will be mirrored from the source to the target.
- **Mirror different files**—Only those protected files that are different based on date and time, size, or attributes will be mirrored from the source to the target.
  - **Mirror only if the file on the source is newer than the copy on the target**—Only those protected files that are newer on the source are mirrored to the target.

    

    If you are using a database application or are protecting a domain controller, do not use this option 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 that might be newer, get mirrored.

  - **Use block checksum for comparisons**—For those files flagged as different, the mirroring process can perform a block checksum comparison and send only those blocks that are different.

File Differences Mirror Options Compared

The following table will help you understand how the various difference mirror options work together, including when you are using the block checksum option configured through the *Source server properties* on page 92.

An X in the table indicates that option is enabled. An X enclosed in parentheses (X) indicates that the option can be on or off without impacting the action performed during the mirror.

Not all job types have the source newer option available.

Source Server Properties	Job Properties			Action Performed
Block Checksum Option	File Differences Option	Source Newer Option	Block Checksum Option	
(X)	X			Any file that is different on the source and target based on the date, time, size, and/or attribute is transmitted to the target. The mirror sends the entire file.
(X)	X	X		Any file that is newer on the source than on the target based on date and/or time is transmitted to the target. The mirror sends the entire file.
	X		X	Any file that is different on the source and target based on date, time, size, and/or attributed is flagged as different. The mirror then performs a checksum comparison on the flagged files and only sends those blocks that are different.
X	X		X	The mirror performs a checksum comparison on all files and only sends those blocks that are different.
(X)	X	X	X	Any file that is newer on the source than on the target based on date and/or time is flagged as different. The mirror then performs a checksum comparison on the flagged files and only sends those blocks that are different.

- **Enable scheduled verification**—Verification is the process of confirming that the source replica data on the target is identical to the original 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, can automatically initiate a remirror, if configured. The remirror ensures data integrity between the source and target. When this option is enabled, Double-Take will verify the source replica data on the target and generate a verification log.

  

  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.

  - **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 at the specified date and time.
  - **Mirror files to the target server when verifying**—When this option is enabled, in addition to verifying the data and generating a log, Double-Take will also mirror to the target any protected files that are different on the source.
    - **Mirror only if the file on the source is newer than the copy on the target**—Only those protected files that are newer on the source are mirrored to the target.

      

      If you are using a database application or are protecting a domain controller, do not use this option 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 that might be newer, get mirrored.

    - **Use block checksum for comparisons**—For those files flagged as different, the mirroring process can perform a block checksum comparison and send only those blocks that are different.
- **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 the calculation is enabled, it is completed before the job starts mirroring, which can take a significant amount of time depending on the number of files and system performance. If your job contains a large number of files, for example, 250,000 or more, you may want to disable the calculation so that data will start being mirrored sooner. 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.

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

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

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 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.

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 95 for more information.

During a mirror, orphaned file processing success messages will be logged to a separate orphaned file log. 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 100 for details on the location of that folder. The orphaned log file is overwritten during each orphaned file processing during a mirror, and the log file will be a maximum of 50 MB.

Target Route

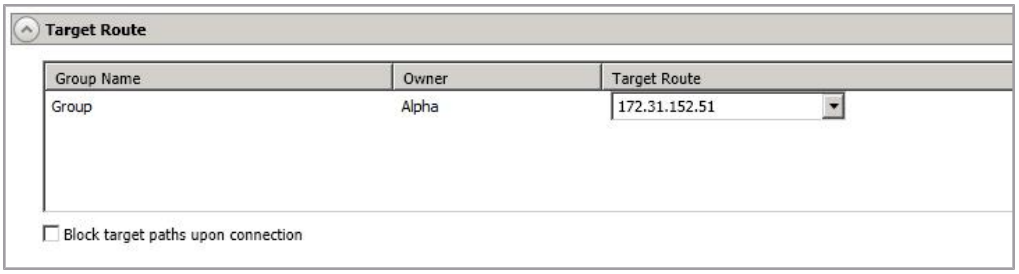

This section is only applicable if your target is a cluster.

- **Target Route**—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 to let Double-Take choose an automatic route.
- **Block target paths upon 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. Any target paths that are blocked will be unblocked automatically during the failover process so that users can modify data after failover.

Network Route

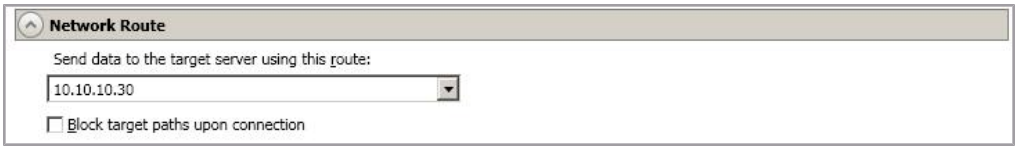

This section is not applicable if your target is a cluster.

- **Send data to the target server using this route**—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 NAT router) if you are using a NAT environment.

  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.

- **Block target paths upon 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. Any target paths that are blocked will be unblocked automatically during the failover process so that users can modify data after failover. During restoration, the paths are automatically blocked again. If you failover and failback without performing a restoration, the target paths will remain unblocked.

Path Mapping

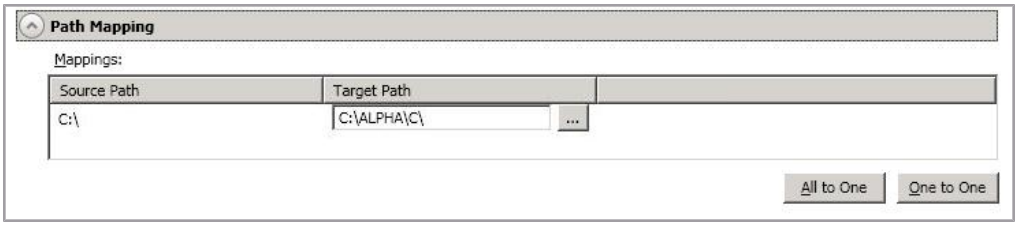

For the **Mappings**, specify the location on the target where the replica of the source data will be stored. By default, the replica source data will be stored in the same directory structure on the target, in a one to one configuration. Make sure you update this location if you are protecting multiple sources or jobs to the same target. If your target is a standalone server, Double-Take offers two pre-defined locations as well as a custom option that allows you to set your path. If your target is a cluster, you only have the custom location.

- **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.
- **Use Defaults**—Click this button to reset the **Target Path** location back to its default settings. This option is only available if your target is a cluster.

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. Keep in mind that this mapping will avoid sharing violations on the target, however during a restoration, you will get sharing violations on the source because the restoration mapping is one to one and your system state files will be in use on the source you are restoring to. In this case, restoration will never complete. If you will need to restore data and you must protect system state data, you should use a full server job.

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 and your location on the target is a non-NTFS 5 volume, 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.

Failover Services

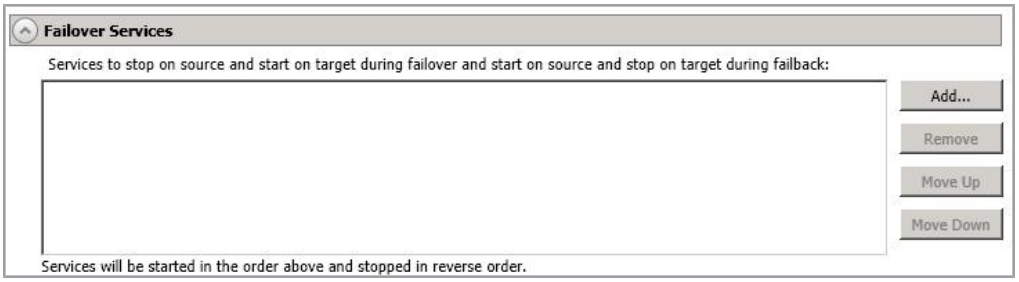

This section is not applicable to clustered environments.

Services to stop on source and start on target during failover and start on source and stop on target during failback—If necessary, you can start and stop services during failover and failback. Click **Add** to insert a service into the list or **Remove** to remove a service from the list. The services will be started in the order they appear and stopped in the reverse order. Highlight a service and click **Move Up** or **Move Down** to arrange the services in the desired order.

Snapshots

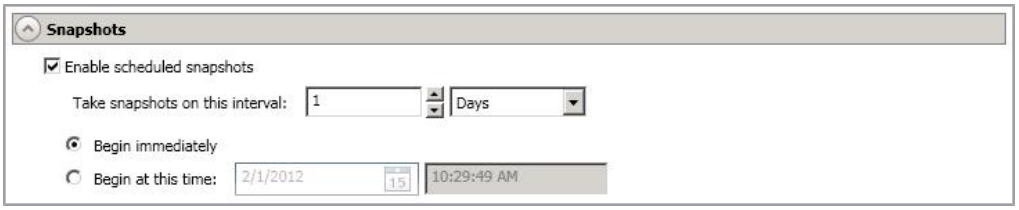

This section is not applicable to clustered environments.

A snapshot is an image of the source replica data on the target taken at a single point in time. You can view the snapshots in VSS and recover any files or folders desired. However, you cannot failover to a snapshot.

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 starting at a later date and time. Specify the date and time parameters to indicate when you want to start.

See *Managing snapshots* on page 161 for details on taking manual snapshots and deleting snapshots.

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

Compression

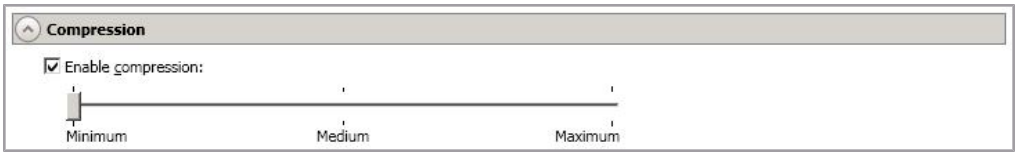

To help reduce the amount of bandwidth needed to transmit Double-Take data, compression allows you to compress data prior to transmitting it across the network. In a WAN environment this provides optimal use of your network resources. If compression is enabled, the data is compressed before it is transmitted 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

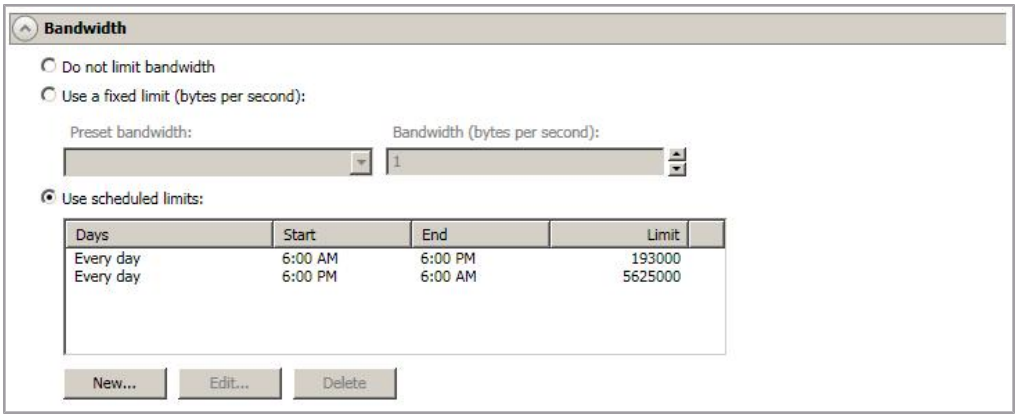

Bandwidth limitations are available to restrict the amount of network bandwidth used for Double-Take data transmissions. When a bandwidth limit is specified, Double-Take never exceeds that allotted amount. The bandwidth not in use by Double-Take is available for all other network traffic.

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

The scheduled option is not available if your source is a cluster.

- **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

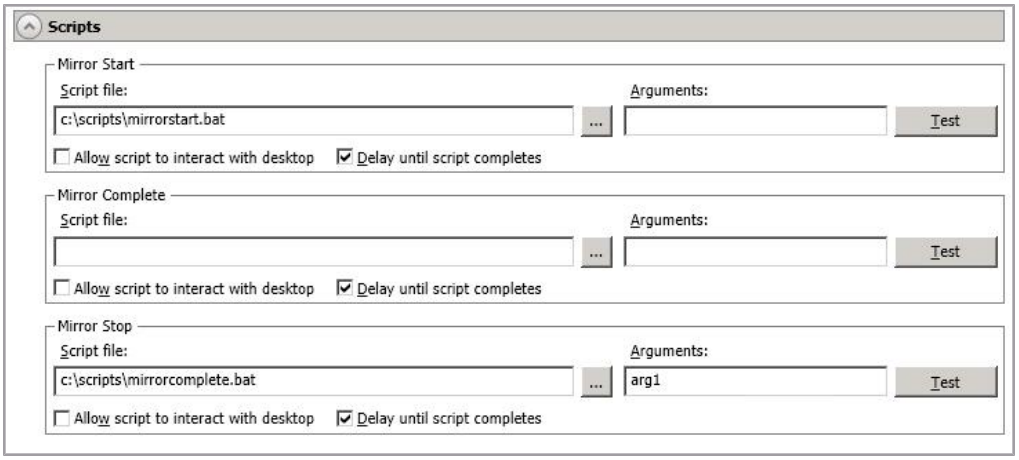

You can customize mirroring by running scripts on the target at pre-defined points in the mirroring process. Scripts may contain any valid Windows command, executable, or batch file. The scripts are processed using the same account running the Double-Take service, unless you have identified a specific account through the server's properties. See *Script credentials* on page 99. 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.
- **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.

Mirror scripts are dependent on the target and the **Target Path Mappings** specified under the **Network Route & Folder Selection** section. 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).

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

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

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

If your target is a cluster, a cluster resource named DTTargetRes_FilesAndFolders_GUID will be created in the target cluster group that contains the disk where data is being replicated from the source. If the target cluster group that contains the DTTargetRes_FilesAndFolders_GUID resource is moved to a different node, an auto-disconnect and reconnect will occur, and a re-mirror will be initiated.

Jobs in a NAT environment may take longer to start.

If you are using a standalone to cluster configuration, the **Failover** button on the **Manage Jobs** page will be enabled once a mirror is complete. Do not failover as unexpected results may occur. Failover is not supported for standalone to cluster configurations.

Managing and controlling files and folders jobs

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

The jobs displayed in the right pane depend on the server group folder selected in the left pane. Every job for each server in your console session is displayed when the **Jobs on All Servers** group is selected. If you have created and populated server groups (see *Managing servers* on page 65), then only the jobs associated with the server or target servers in that server group will be displayed in the right pane.

- See *Overview job information displayed in the top pane* on page 207
- See *Detailed job information displayed in the bottom pane* on page 209
- See *Job controls* on page 211

Overview job information displayed in the top pane

The top pane displays high-level overview information about your jobs.

Column 1 (Blank)

The first blank column indicates the state of the job.

 The job is in a healthy state.

 The job is in a warning state. This icon is also displayed on any server groups that you have created that contain a job in a warning state.

 The job is in an error state. This icon is also displayed on any server groups that you have created that contain a job in an error state.

 The job is in an unknown state.

Job

The name of the job

Source Server

The name of the source. This could be a name or IP address of a standalone server, a cluster, or a node. Cluster jobs will be associated with the cluster name and standalone jobs will be associated with a standalone server or a cluster node.

Target Server

The name of the target. This could be a name or IP address of a standalone server, a cluster, or a node. Cluster jobs will be associated with the cluster name and standalone jobs will be associated with a standalone server or a cluster node.

Job Type

Each job type has a unique job type name. This job is a Files and Folders job. For a complete list of all job type names, press F1 to view the Double-Take Console online help.

Activity

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

Mirror Status

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

Replication Status

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

Transmit Mode

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

Detailed job information displayed in the bottom pane

The details displayed in the bottom pane of the **Manage Jobs** page provide additional information for the job highlighted in the top pane. If you select multiple jobs, the details for the first selected job will be displayed.

Name

The name of the job

Target data state

- **OK**—The data on the target is in a good state.
- **Mirroring**—The target is in the middle of a mirror process. The data will not be in a good state until the mirror is complete.
- **Mirror Required**—The data on the target is not in a good state because a remirror is required. This may be caused by an incomplete or stopped mirror or an operation may have been dropped on the target.
- **Restore Required**—The data on the source and target do not match because of a failover condition. Restore the data from the target back to the source. If you want to discard the changes on the target, you can remirror to resynchronize the source and target.
- **Snapshot Reverted**—The data on the source and target do not match because a snapshot has been applied on the target. Restore the data from the target back to the source. If you want to discard the changes on the target, you can remirror to resynchronize the source and target.
- **Busy**—The source is low on memory causing a delay in getting the state of the data on the target.
- **Not Loaded**—Double-Take target functionality is not loaded on the target server. This may be caused by an activation code error.
- **Unknown**—The console cannot determine the status.

Mirror remaining

The total number of mirror bytes that are remaining to be sent from the source to the target

Mirror skipped

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

Replication queue

The total number of replication bytes in the source queue

Disk queue

The amount of disk space being used to 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**.

Connected since

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

Recent activity

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

Additional information

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

Job controls

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

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

Create a New Job 

This button leaves the **Manage Jobs** page and opens the **Get Started** page.

View Job Details 

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

Delete 

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 77. You will remain on the **Manage Jobs** page after updating the server credentials. If your servers use the same credentials, make sure you also update the credentials on the **Manage Servers** page so that the Double-Take Console can authenticate to the servers in the console session. See *Managing servers* on page 65.

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. Failover monitoring will continue 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. Failover monitoring will continue while the job is stopped. 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.

Snapshots are not applicable to clustered environments.

Manage Snapshots 

Allows you to manage your snapshots by taking and deleting snapshots for the selected job. See *Managing snapshots* on page 161 for more information.

Snapshots are not applicable to clustered environments.

Failover or Cutover 

Starts the failover process. See *Failing over files and folders jobs* on page 225 for the process and details of failing over a files and folders job.

Failback 

Starts the failback process. See *Failback and restoration for files and folders jobs* on page 226 for the process and details of failing back a files and folders job.

Restore 

Starts the restoration process. See *Failback and restoration for files and folders jobs* on page 226 for the process and details of restoring a files and folders job.

Reverse 

Reverses protection. Reverse protection does not apply to files and folders jobs.

Undo Failover 

Cancels a test failover by undoing it. Undo failover does not apply to files and folders jobs.

View Job Log 

Opens the job log. On the right-click menu, this option is called **View Logs**, and you have the option of opening the job log, source server log, or target server log. See *Viewing the log files through the Double-Take Console* on page 678 for details on all three of these logs.

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 all files**—All protected files will be mirrored from the source to the target.
  - **Mirror different files**—Only those protected files that are different based on date and time, size, or attributes will be mirrored from the source to the target.
    - **Mirror only if the file on the source is newer than the copy on the target**—Only those protected files that are newer on the source are mirrored to the target.

    

    If you are using a database application or are protecting a domain controller, do not use this option 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 that might be newer, get mirrored.

- **Use block checksum for comparisons**—For those files flagged as different, the mirroring process can perform a block checksum comparison and send only those blocks that are different.

- **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 the calculation is enabled, it is completed before the job starts mirroring, which can take a significant amount of time depending on the number of files and system performance. If your job contains a large number of files, for example, 250,000 or more, you may want to disable the calculation so that data will start being mirrored sooner. 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.
  - **Create verification report only**—This option verifies the data and generates a verification log, but it does not remirror any files that are different on the source and target. See *Verification log* on page 102 for details on the log file.
  - **Mirror files to the target server automatically**—When this option is enabled, in addition to verifying the data and generating a log, Double-Take will also mirror to the target any protected files that are different on the source.
    - **Mirror only if the file on the source is newer than the copy on the target**—Only those protected files that are newer on the source are mirrored to the target.

      

      If you are using a database application or are protecting a domain controller, do not use this option 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 that might be newer, get mirrored.

    - **Use block checksum for comparisons**—For those files flagged as different, the mirroring process can perform a block checksum comparison and send only those blocks that are different.
- **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.
- **Update Shares**—Windows share information is automatically updated on the target once an hour. This option allows you to manually update share information immediately when the option is selected. Shares are not applicable to environments where the target is a cluster.

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 65.

Type a server name

Displays only jobs that contain the text you entered. If you have created and populated server groups, then only jobs that contain the text you entered associated with the server or target servers in that server group will be displayed. See *Managing servers* on page 65.

Overflow Chevron 

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

Viewing files and folders job details

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

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

Job name

The name of the job

Job type

Each job type has a unique job type name. This job is a Files and Folders job. For a complete list of all job type names, press F1 to view the Double-Take Console online help.

Health

 The job is in a healthy state.

 The job is in a warning state.

 The job is in an error state.

 The job is in an unknown state.

Activity

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

Connection ID

The incremental counter used to number connections. The number is incremented when a connection is created. It is also incremented by internal actions, such as an auto-disconnect and auto-reconnect. The lowest available number (as connections are created, stopped, deleted, and so on) will always be used. The counter is reset to one each time the Double-Take service is restarted.

Transmit mode

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

Target data state

- **OK**—The data on the target is in a good state.
- **Mirroring**—The target is in the middle of a mirror process. The data will not be in a good state until the mirror is complete.
- **Mirror Required**—The data on the target is not in a good state because a remirror is required. This may be caused by an incomplete or stopped mirror or an operation may have been dropped on the target.
- **Restore Required**—The data on the source and target do not match because of a failover condition. Restore the data from the target back to the source. If you want to discard the changes on the target, you can remirror to resynchronize the source and target.
- **Snapshot Reverted**—The data on the source and target do not match because a snapshot has been applied on the target. Restore the data from the target back to the source. If you want to discard the changes on the target, you can remirror to resynchronize the source and target.
- **Busy**—The source is low on memory causing a delay in getting the state of the data on the target.
- **Not Loaded**—Double-Take target functionality is not loaded on the target server. This may be caused by an activation code error.
- **Unknown**—The console cannot determine the status.

Target route

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

Compression

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

Encryption

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

Bandwidth limit

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

Connected since

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

Additional information

Depending on the current state of your job, you may see additional information displayed to keep you informed about the progress and status of your job. If there is no

additional information, you will see (None) displayed.

Mirror status

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

Mirror percent complete

The percentage of the mirror that has been completed

Mirror remaining

The total number of mirror bytes that are remaining to be sent from the source to the target

Mirror skipped

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

Replication status

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

Replication queue

The total number of replication bytes in the source queue

Disk queue

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

Bytes sent

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

Bytes sent compressed

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

Validating a files and folders job

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

1. From the **Manage Jobs** page, highlight the job and click **View Job Details** in the toolbar.
2. In the **Tasks** area on the right on the **View Job Details** page, click **Validate job properties**.
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 enable protection. Depending on the error, you may be able to click **Fix** or **Fix All** and let Double-Take correct the problem for you. For those errors that Double-Take cannot correct automatically, you will need to modify the source or target to correct the error, or you can select a different target. You must revalidate the selected servers, by clicking **Recheck**, until the validation check passes without errors.

   Validation checks for an existing job are logged to the job log on the target server. See *Log files* on page 677 for details on the various log files.
4. Once your servers have passed validation, click **Close**.

Editing a files and folders job

Use these instructions to edit a files and folders job.

1. From the **Manage Jobs** page, highlight the job and click **View Job Details** in the toolbar.
2. In the **Tasks** area on the right on the **View Job Details** page, click **Edit job properties**. (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 have the same options available for your files and folders job as when you created the job. See *Creating a files and folders job* on page 176 for details on each job option.

   

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

4. If you want to modify the workload items or replication rules for the job, click **Edit workload or replication rules**. Modify the **Workload item** you are protecting, if desired. Additionally, you can modify the specific **Replication Rules** for your job.

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

   You can also enter wildcard rules, however you should do so carefully. Rules are applied to files that are closest in the directory tree to them. If you have rules that include multiple folders, an exclusion rule with a wild card will need to be added for each folder that it needs applied to. For example, if you want to exclude all .log files from D:\ and your rules include D:\, D:\Dir1 , and D:\Dir2, you would need to add the exclusion rule for the root and each subfolder rule. So you will need to add exclude rules for D:*.log , D:\Dir1*.log, and D:\Dir2*.log.

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

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

   Errors are designated by a white X inside a red circle. Warnings are designated by a black exclamation point (!) inside a yellow triangle. A successful validation is designated by a white checkmark inside a green circle. You can sort the list by the icon to see errors, warnings, or successful validations together. Click on any of the validation items to see details. You must correct any errors before you can enable protection. Depending on the error, you may be able to

click **Fix** or **Fix All** and let Double-Take correct the problem for you. For those errors that Double-Take cannot correct automatically, you will need to modify the source or target to correct the error, or you can select a different target. You must revalidate the selected servers, by clicking **Recheck**, until the validation check passes without errors.

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.

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

Viewing a files and folders job log

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

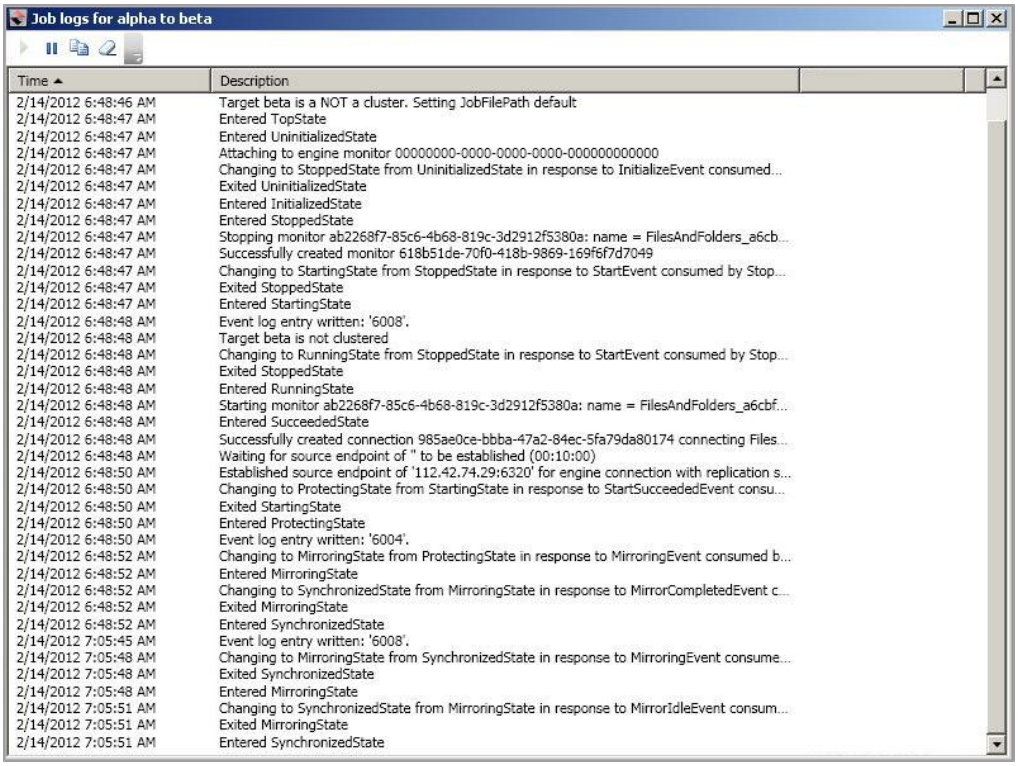

The following table identifies the controls and the table columns in the **Job logs** window.

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.

Failing over files and folders jobs

When a failover condition has been met, failover will be triggered automatically if you disabled the wait for user option during your failover configuration. If the wait for user before failover option is enabled, you will be notified in the console when a failover condition has been met. At that time, you will need to trigger it manually from the console when you are ready.

If you are using a files and folders job in a standalone to cluster configuration, the **Failover** button will be enabled once a mirror is complete. Do not failover as unexpected results may occur. Failover is not supported for files and folders jobs in a standalone to cluster configuration.

1. On the **Manage Jobs** page, highlight the job that you want to failover and click **Failover or Cutover** in the toolbar.
2. Select the type of failover to perform.
   - **Failover to live data**—Select this option to initiate a full, live failover using the current data on the target. The target will stand in for the source by assuming the network identity of the failed source. User and application requests destined for the source server or its IP addresses are routed to the target.
   - **Perform test failover**—This option is not available for files and folders jobs.
   - **Failover to a snapshot**—This option is not applicable to files and folders jobs.
3. Select how you want to handle the data in the target queue. You may want to check the amount of data in queue on the target by reviewing the *Statistics* on page 688 or *Performance Monitor* on page 790.
   - **Apply data in target queues before failover or cutover**—All of the data in the target queue will be applied before failover begins. The advantage to this option is that all of the data that the target has received will be applied before failover begins. 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 failover or cutover immediately**—All of the data in the target queue will be discarded and failover will begin immediately. The advantage to this option is that failover will occur immediately. The disadvantage is that any data in the target queue will be lost.
4. When you are ready to begin failover, click **Failover**.

IPv6 addresses on the source will be set to DHCP on the target after failover. Update them to static addresses manually, if needed.

If your NICs were configured for network load balancing (NLB), you will have to reconfigure that after failover.

Failback and restoration for files and folders jobs

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

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

Ideally, you want to restore your data from the target back to the source before you failback. This allows users who are currently accessing their data on the target because of failover to continue accessing their data. Restoration before failback reduces user downtime. Another method, which may be easier in some environments that have strict IP addressing policies, allows you to failback first and then restore the data from the target to the source. A possible disadvantage to this process is that users may experience longer downtime, depending on the amount of data to be restored, because they will be unable to access their data during both the restoration and the failback.

- See *Restoring then failing back files and folders jobs* on page 227
- See *Failing back then restoring files and folders jobs* on page 230

If you want to failback only, without performing a restoration, follow the instructions for failing back then restoring and you will be able to skip the restoration process. Keep in mind that if you skip the restoration process, any data changes that were made on the target during failover will be lost.

Restoring then failing back files and folders jobs

Restoring before failing back allows your users to continue accessing their data on the failed over target, which is standing in for the source, while you perform the restoration process. The key to this process is to keep the users off of the source, but allow the source and target to communicate to perform the restoration.

1. Locate the file connect.sts on the source where you installed Double-Take and rename it to connect.sts.old. This will keep the original connection from attempting to reconnect when you bring the source online.
2. Resolve the problem(s) on the source that caused it to fail. Make sure in resolving the problems, that you do not bring the source on the network at this time because the target currently has the source's identity because of failover.
3. Disable all of the IP addresses on the source that you failed over to the target.
4. If you failed over all of your source IP addresses, change an existing IP address on the source to a new, unique IP address that the target can access. If you inserted your source server into the console using a private IP address when you created the job, and you did not failover that private IP address, you can skip this step.
5. Configure your new, unique IP address that you created or the private IP address that you are using so that it does not automatically register with DNS. This option is on the Advanced TCP/IP Settings dialog box on the DNS tab.
6. Bring the source onto the network using the IP address that the target can access. This will either be the new, unique IP address or the private IP address. You can disregard any identity conflict errors.
7. Stop any applications that may be running on the source. Files must be closed on the source so that updated files from the target will overwrite the files on the source.
8. If you had to create a new, unique IP address on the source, you will have to remove the original source in the console and add it back in using the new, unique IP address. Use a local account, not a domain account, that is a member of the Double-Take Admin and Administrators groups. Complete this step on the **Manage Servers** page. If you inserted your source server into the console using a private IP address when you created the job, you can skip this step.
9. On the **Manage Jobs** page, highlight the job and click **Restore**.
10. Confirm **Restore data** is selected, then highlight your source server in the server list. Look for the new, unique IP address or the private IP address in parenthesis after the server name.

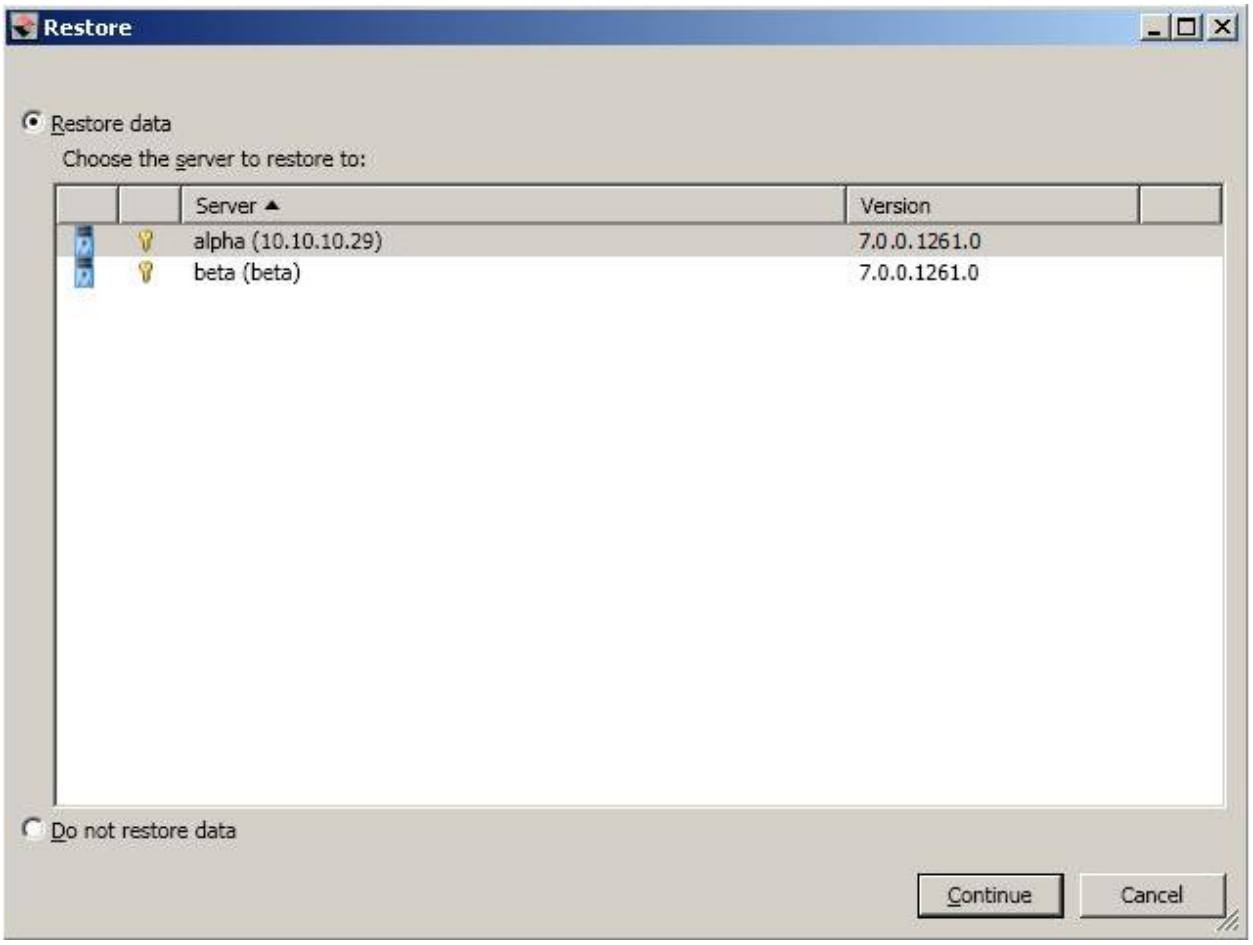

11. Click **Continue** to start the restoration.

During the restoration, only the data is restored back to the source. Shares are not created on the source during the restoration. Shares that were created on the target during failover will need to be created manually on the source, if they do not already exist.

12. During the restoration process, the **Activity** and **Mirror Status** will both indicate **Restoring**. When the restoration is complete, the **Mirror Status** will change to **Idle** and the **Activity** will change to **Restored**.
13. On the **Manage Jobs** page, highlight the job and click **Failback**.
14. In the dialog box, highlight the job that you want to failback and click **Failback**.
15. After failback is complete, stop the job by clicking **Stop**.

If you do not see your job after failback, remove and re-add your target to the console. The job may not be visible depending on where you are running the console from.

16. After the job is stopped, enable or add the IP addresses on the source that you disabled earlier. Make sure that you keep any new addresses that you created because the job is still using that address.
17. If you restored to a new source and are going to enable protection again, edit the job to reconfigure your failover settings.
18. Click **Start** to restart protection.

Failing back then restoring files and folders jobs

Failback before restoration can be a simpler process, but it may require additional downtime. The amount of downtime will depend on the amount of data to be restored. Users must be kept off of the source and target during this entire process.

1. Remove the source from the network and fix the issue that caused your source server to fail. Make sure in resolving the problems that you do not bring the source on the network at this time because the target currently has the source's identity because of failover.
2. Schedule a time for the failback and restoration process. Select a time that will have minimal disruption on your users.
3. When you are ready to begin the failback process, prohibit user access to both the source and target.
4. On the **Manage Jobs** page, highlight the job and click **Failback**.
5. Highlight the job that you want to failback and click **Failback**.
6. Once failback is complete, the **Activity** will indicate **Restore required**. At this time, bring the source onto the network. Make sure that end users continue to be prohibited from accessing both the source and target because the updated data from the target needs to be restored back to the source.

   

   Depending on where you are running the console from, you may need to add the target back to the console after failback in order to see your job.

7. Stop any applications that may be running on the source. Files must be closed on the source so that updated files from the target will overwrite the files on the source.
8. On the **Manage Jobs** page, highlight the job and click **Restore**.

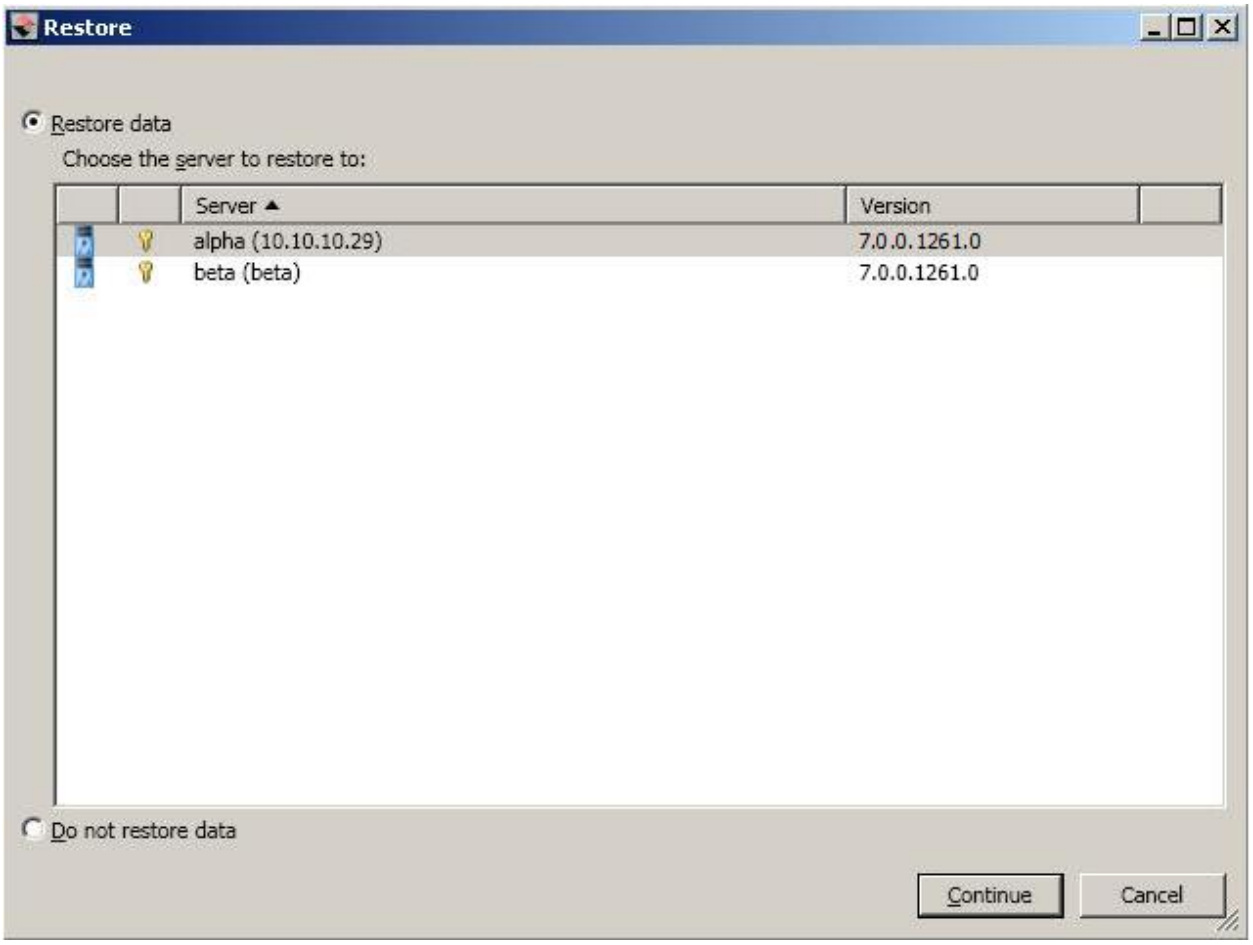

9. If you want to skip the restoration, select **Do not restore data**, and click **Continue**. Keep in mind that if you skip the restoration process, any data changes that were made on the target during failover will be lost. If you want to restore the changed data from the target back to the source, select **Restore data**, highlight your source server in the server list, and click **Continue** to start the restoration.

During the restoration, only the data is restored back to the source. Shares are not created on the source during the restoration. Shares that were created on the target during failover will need to be created manually on the source, if they do not already exist.

---

10. When the restoration is complete, the **Activity** will indicate **Mirror required**. If you skipped the restoration, the **Activity** will also indicate **Mirror required**. At this point, stop the job by clicking **Stop**.
11. After the job is stopped, you can allow users to access the source again.
12. If you restored to a new source and are going to enable protection again, edit the job to reconfigure your failover settings.
13. Click **Start** to restart protection.

Chapter 8 Full server protection

Create a full server job when you want to protect the entire source, including the server's system state. You can also use it to protect an application server. This type of job is the most flexible, allowing you to go from physical to physical, physical to virtual, virtual to virtual, and virtual to physical.

- See *Full server requirements* on page 233—Full server protection includes specific requirements for this type of protection.
- See *Creating a full server job* on page 238—This section includes step-by-step instructions for creating a full server job.
- See *Managing and controlling full server jobs* on page 264—You can view status information about your full server jobs and learn how to control these jobs.
- See *Failing over full server jobs* on page 282—Use this section when a failover condition has been met or if you want to failover manually.
- See *Reversing full server jobs* on page 285—Use this section to reverse protection. The source (what was your original target hardware) is now sending data to the target (what was your original source hardware).

Full server requirements

In addition to the *Core Double-Take requirements* on page 23, use these requirements for full server protection.

- **Operating system**—The source and target servers can be any operating system listed in the *Core Double-Take requirements* on page 23. However, keep in mind that a target server may meet these requirements but may not be suitable to stand-in for a source in the event of a source failure. See *Target compatibility* on page 236 for additional information regarding an appropriate target server for your particular source.
- **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 failover 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.
- **Source storage**—If you will be enabling reverse protection, the source must have enough space to store, process, and apply the target's system state date. See the disk space requirements in *Target compatibility* on page 236 for details on the space requirements for various operating systems.
- **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 failing over an image of a server to different hardware.
- **Microsoft Server Core 2008 R2, 2012, or 2012 R2**—These operating systems are only supported in a Server Core to Server Core configuration. Additionally, DNS updates are not supported for Server Core servers.
- **Hyper-V servers**—Full server protection of a Hyper-V server is not supported.
- **Network adapters**—NIC teaming is not supported for full server protection.
- **NAT support**—Full server jobs can support NAT environments only when reverse protection is disabled. Additionally, your NAT environment must be an IP-forwarding configuration with one to one port mappings. Port-forwarding is not supported. Also, only IPv4 is supported for NAT environments. Make sure you have added your servers to the Double-Take Console using the correct IP address. Review the *NAT configuration* table on page 72 in the *Adding servers* section before you start the job creation process.
- **Microsoft .NET Framework**—Microsoft .NET Framework version 3.5 Service Pack 1 is required on the source and target. This version is not included in the .NET version 4.0 release. Therefore, even if you have .NET version 4.0 installed, you will also need version 3.5.1. If you are using Windows 2008 or earlier, you can install this version from the Double-Take DVD, via a web connection during the Double-Take installation, or from a copy you have obtained manually from the Microsoft web site.If you are using Windows 2008 R2 or later, you can enable it through Windows features.
- **Snapshots**—You can take and failover to Double-Take snapshots using a full server job. See *Core Double-Take requirements* on page 23 for the specific snapshot requirements.

- **Supported configurations**—The following table identifies the supported configurations for a full server job.

Configuration		Supported	Not Supported
Source to target configuration	One to one, active/standby	X	
	One to one, active/active		X
	Many to one		X
	One to many [1]	X	
	Chained		X
	Single server		X
Server configuration	Standalone to standalone	X	
	Standalone to cluster		X
	Cluster to standalone [2]		X
	Cluster to cluster		X
	Cluster Shared Volumes (CSV) guest level	X	
	Cluster Shared Volumes (CSV) host level		X
Upgrade configuration [3]	Upgrade 5.3 full server job to 7.0 full server job	X	
	Upgrade 6.0 full server job to 7.0 full server job	X	
Version 7.0 console[4]	Version 7.0 console can create job for 5.3 source and 5.3 target	X	
	Version 7.0 console can create job for 6.0 source and 6.0 target	X	
	Version 7.0 console can create job for 7.0 source and 7.0 target	X	

1. If you are using a one to many configuration, you will only be able to configure reverse protection for the first job. Subsequent jobs from that source will have reverse protection disabled.
2. Full server protection of a cluster is not a default supported configuration, but it is possible with assistance from Professional Services. If you want to use a full server job in a cluster environment, contact Sales or Professional Services. Because this is an advanced configuration, you will be referred to Professional Services for further assistance.
3. When upgrading, you can perform a rolling upgrade where you update the target server first. After the upgrade is complete, the source will automatically reconnect to the target.

Upgrade the source when convenient. However, after failover you will not be able to reverse until the original source has been upgraded.

4. Newer job options available in the version 7.0 console will not be functional when creating jobs for servers running version 5.3 or 6.0.

Target compatibility

- **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 2003 and 2003 R2 are considered the same operating system, however the Windows 2008 and 2012 and their R2 releases are considered different operating systems. Therefore, you can have Windows 2003 on the source and Windows 2003 R2 on the target, but 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.
- **Windows Azure**—Because Windows Azure uses remote desktop (RDP) for console connections to virtual machines running on Azure, if your target is running on Windows Azure, you must have remote desktop enabled on the source or the target will be inaccessible after failover.
- **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.

  Additionally, if your source is a domain controller, you will not be able to reverse protection.
- **Architecture**—The source and the target must have the same architecture. For example, you cannot failover a 32-bit server to a 64-bit server.
- **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 same 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, however you will be unable to reverse protection if the target has more drives than the source.
- **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.
- **Double-Take version**—If you will be using the reverse feature with your full server job, your source and target must be running the same Double-Take version.
- **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. If you will be enabling reverse protection, the source must have enough space to store, process, and apply the target's system state data. In either case, 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 source 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 this 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 2003**—at least 2-3 GB
  - **Windows 2008**—at least 7-9 GB
  - **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.

Creating a full server job

Use these instructions to create a full server job.

If you are creating a full server job in order to revert back to your original configuration after failing over a full server to ESX or full server to Hyper-V, you will need to perform a few additional tasks before creating the full server job. Contact technical support if you need assistance with these steps.

1. On either the source or target, stop the Double-Take and Double-Take Management Service services.
2. Remove the GUID value from HKEY_LOCAL_MACHINE\SOFTWARE\NSI Software\Double-Take\CurrentVersion\Communication\UniqueId. Do not delete the UniqueId key. Only delete the GUI value within the key.
3. Restart the the Double-Take and Double-Take Management Service services.
4. Remove and then add your servers back into the Double-Take Console.
5. Install a different license on the original source and complete a host transfer if necessary.

---

1. Review these best practices before you create your job.
   - **NIC configuration**—If you are planning to failover the IP address of the source, use a separate NIC and separate network for a Double-Take reserved IP address that will not be failed over. If you are unable to do that and just one NIC is used for both production and reserved IP addresses, disable DNS registration on the NIC. If you are not going to failover the IP address of the source, an additional NIC and address is not necessary. In this case, Double-Take will block the DNS record for that address while it is failed over.
   - **Windows 2003 servers**—If your servers are running Windows 2003 and you have a separate, reserved NIC configuration, make sure that the reserved NIC address does not have a default gateway set. If your servers are running Windows 2003 and you have a single NIC configuration, make the Double-Take reserved IP address is the primary IP address in the NIC properties on the source, otherwise, the failed over IP address may need to be removed on the original source after failover to allow the reserved IP address to come online.
   - **Third machine setup**—Ideally, you should use a third machine (not the source or target) for configure protection, to failover, and to reverse. The third machine must be able to communicate with the source and target using their reserved IP addresses.
   - **Double-Take Console**—Insert your source and target servers into the console using the reserved IP addresses and a local computer account that is a member of the Double-Take Admin and Administrators groups.
   - **NAT environments**—Full server jobs in a NAT environment are only supported with reverse protection disabled. Additionally, you must make sure you have added your servers to the Double-Take Console using the correct IP address. Review the *NAT configuration* table on page 72 in the *Adding servers* section before you start the job creation process.
2. Click **Get Started** from the toolbar.

3. Select **Double-Take Availability** and click **Next**.
4. Select **Protect files and folders, an application, or an entire Windows server** and click **Next**.
5. Choose your source server. This is the physical or virtual server that you want to protect.

- **Current Servers**—This list contains the servers currently available in your console session. Servers that are not licensed for the workflow you have selected will be filtered out of the list. Select your source server from the list.
- **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 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.

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. If the source is a domain controller, Double-Take will add the security groups to the users OU, therefore permissions must be located there. If your source is the only domain controller in your network, the account must also be a local account in the local administrators group on the target. If you want Double-Take to update DNS during failover, the account must be a member of the Domain Admins group. If your security policies do not allow use of this group, see *DNS* on page 816 and use the instructions under the Double-Take DFO utility to use a non-Domain Admins account.

6. Click **Next** to continue.
7. Choose the type of workload that you want to protect. Under **Server Workloads**, in the **Workload types** pane, select **Full 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.

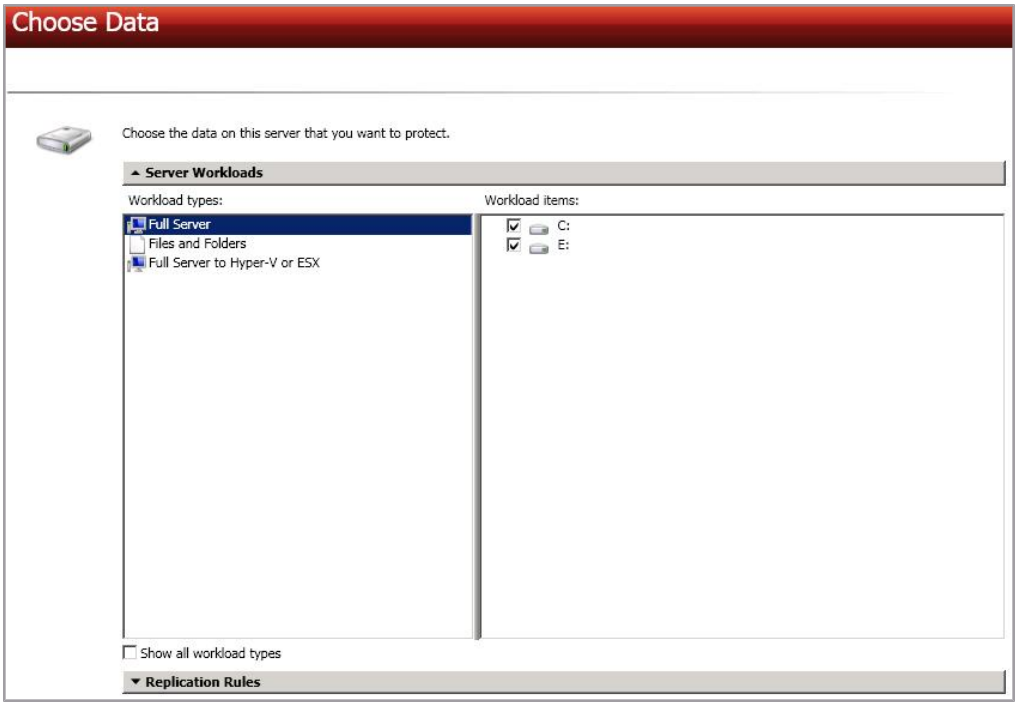

8. By default, Double-Take selects your entire source for protection. If desired, click the **Replication Rules** heading and expand the volumes under **Folders**. You will see that Double-Take automatically excludes particular files that cannot be used during the protection. If desired, you can exclude other files that you do not want to protect, but be careful when excluding data. Excluded volumes, folders, and/or files may compromise the integrity of your installed applications. There are some volumes, folders, and files (identified in italics text) that you will be unable to exclude, because they are required for protection. For example, the boot files cannot be excluded because that is where the system state information is stored.

   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.

If IIS is being used as a hardware platform manager by your hardware vendor on both the source and target, you need to remove the INetPub directory from replication under the **Replication Rules** heading. If IIS is being used as a software application on your source but as a hardware platform manager by your hardware vendor on your target, you need to add the INetPub directory to the **Staged Folders Options** list on the **Set Options** page later in this workflow.

---

9. Click **Next** to continue.
10. Choose your target server. This is the server that will store the replica data from the source.

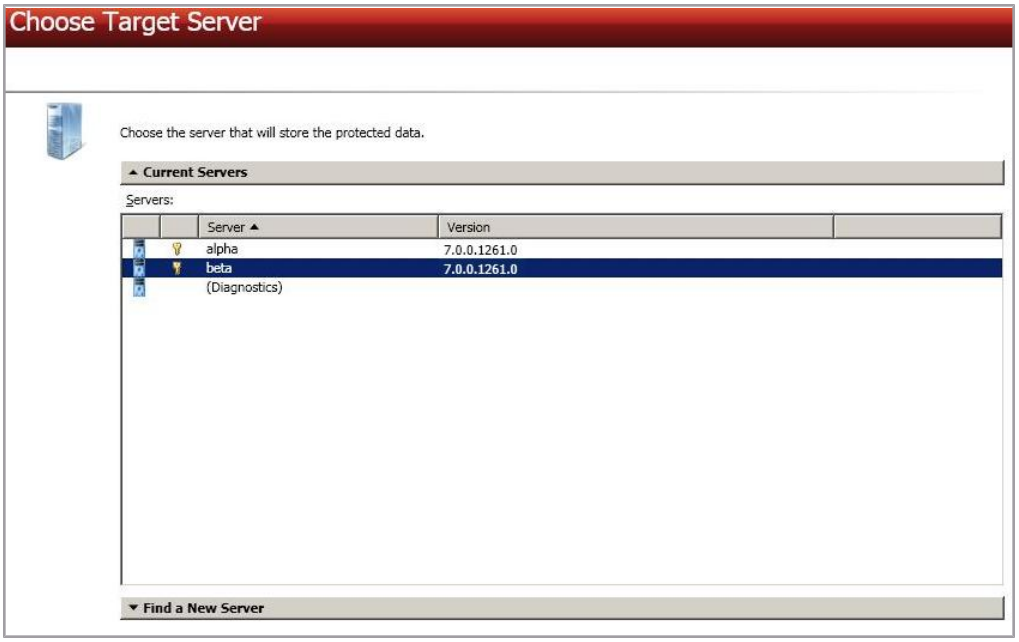

- **Current Servers**—This list contains the servers currently available in your console session. Servers that are not licensed for the workflow you have selected and those not

applicable to the workload type you have selected will be filtered out of the list. Select your target server from the list.

- **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. If the source is a domain controller, Double-Take will add the security groups to the users OU, therefore permissions must be located there. If your source is the only domain controller in your network, the account must also be a local account in the local administrators group on the target. If you want Double-Take to update DNS during failover, the account must be a member of the Domain Admins group. If your security policies do not allow use of this group, see *DNS* on page 816 and use the instructions under the Double-Take DFO utility to use a non-Domain Admins account.

11. Click **Next** to continue.
12. You have many options available for your full server 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 263.

    - *General* on page 243
    - *Failover Monitor* on page 244
    - *Failover Options* on page 246
    - *Failover Identity* on page 247
    - *Creating a full server job* on page 238
    - *Network Adapter Options* on page 251
    - *Mirror, Verify & Orphaned Files* on page 252
    - *Staging Folder Options* on page 256
    - *Target Services* on page 257
    - *Snapshots* on page 258
    - *Compression* on page 259
    - *Bandwidth* on page 260
    - *Scripts* on page 262

General

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

Failover Monitor

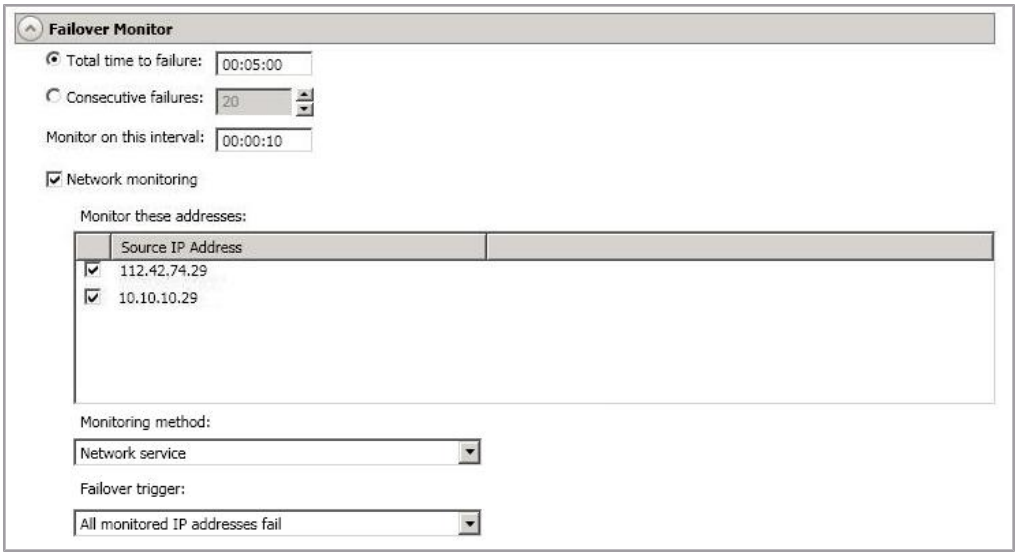

- **Total time to failure**—Specify, in hours:minutes:seconds, how long the target will keep trying to contact the source before the source is considered failed. This time is precise. If the total time has expired without a successful response from the source, this will be considered a failure.

  Consider a shorter amount of time for servers, such as a web server or order processing database, which must remain available and responsive at all times. Shorter times should be used where redundant interfaces and high-speed, reliable network links are available to prevent the false detection of failure. If the hardware does not support reliable communications, shorter times can lead to premature failover. Consider a longer amount of time for machines on slower networks or on a server that is not transaction critical. For example, failover would not be necessary in the case of a server restart.

- **Consecutive failures**—Specify how many attempts the target will make to contact the source before the source is considered failed. For example, if you have this option set to 20, and your source fails to respond to the target 20 times in a row , this will be considered a failure.
- **Monitor on this interval**—Specify, in hours:minutes:seconds, how long to wait between attempts to contact the source to confirm it is online. This means that after a response (success or failure) is received from the source, Double-Take will wait the specified interval time before contacting the source again. If you set the interval to 00:00:00, then a new check will be initiated immediately after the response is received.

  If you choose **Total time to failure**, do not specify a longer interval than failure time or your server will be considered failed during the interval period.

  If you choose **Consecutive failures**, your failure time is calculated by the length of time it takes your source to respond plus the interval time between each response, times the number of consecutive failures that can be allowed. That would be (response time + interval) * failure number. Keep in mind that timeouts from a failed check are included in the response time, so your failure time will not be precise.

- **Network monitoring**—With this option, the target will monitor the source using a network ping.
  - **Monitor these addresses**—Select each **Source IP Address** that you want the target to monitor. If you want to monitor additional addresses, enter the address and click **Add**. In a NAT environment, you can add any additional public IP addresses on the source that may not be listed. Additionally, you should not monitor any private IP addresses on the source because the target cannot reach the source's private address in a NAT environment, thus causing an immediate failure.
  - **Monitoring method**—This option determines the type of network ping used for failover monitoring.
    - **Network service**—Source availability will be tested by an ICMP ping to confirm the route is active.
    - **Replication service**—Source availability will be tested by a UDP ping to confirm the Double-Take service is active.
    - **Network and replication services**—Source availability will be tested by both an ICMP ping to confirm the route is active and a UDP ping to confirm the Double-Take service is active. Both pings must fail in order to trigger a failover.
  - **Failover trigger**—If you are monitoring multiple IP addresses, specify when you want a failover condition to be triggered.
    - **One monitored IP address fails**—A failover condition will be triggered when any one of the monitored IP addresses fails. If each IP address is on a different subnet, you may want to trigger failover after one fails.
    - **All monitored IP addresses fail**—A failover condition will be triggered when all monitored IP addresses fail. If there are multiple, redundant paths to a server, losing one probably means an isolated network problem and you should wait for all IP addresses to fail.

Failover Options

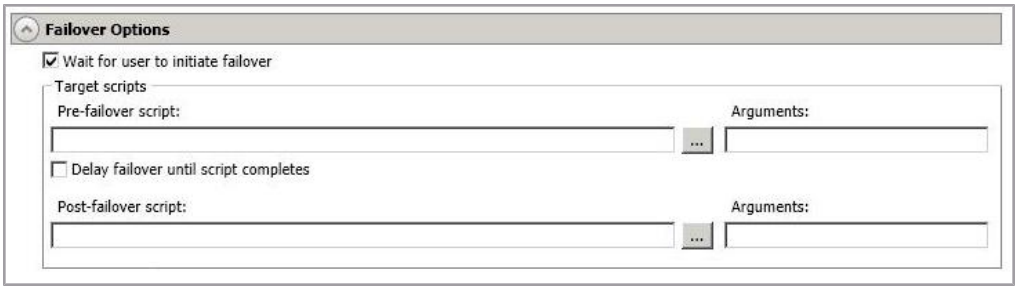

- **Wait for user to initiate failover**—By default, the failover process will wait for you to initiate it, allowing you to control when failover occurs. When a failure occurs, the job will wait in **Failover Condition Met** for you to manually initiate the failover process. Disable this option only if you want failover to occur immediately when a failure occurs.
- **Scripts**—You can customize failover by running scripts on the target or the recovered source. 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 99. Examples of functions specified in scripts include stopping services on the target before failover because they may not be necessary while the target is standing in for the source, stopping services on the target that need to be restarted with the source's machine name and/or IP address, starting services or loading applications that are in an idle, standby mode waiting for failover to occur, notifying the administrator before and after failover occurs, and so on. There are two types of failover scripts.
  - **Pre-failover script**—This script runs on the target at the beginning of the failover process. Specify the full path and name of the script file.
  - **Post-failover script**—This script runs on the recovered source at the end of the failover process. Specify the full path and name of the script file.
  - **Arguments**—Specify a comma-separated list of valid arguments required to execute the script.
  - **Delay until script completes**—Enable this option if you want to delay the failover process until the associated script has completed. If you select this option, make sure your script handles errors, otherwise the failover process may never complete if the process is waiting on a script that cannot complete.

  Scripts will run but will not be displayed on the screen if the Double-Take service is not set to interact with the desktop. Enable this option through the Windows Services applet.

Failover Identity

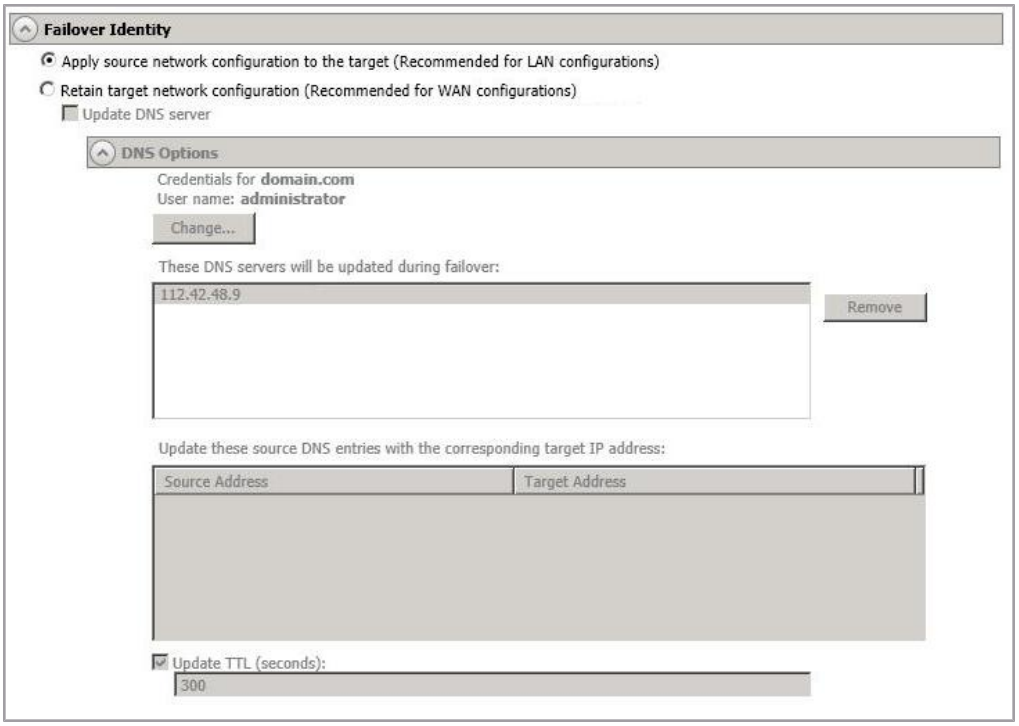

- **Apply source network configuration to the target**—If you select this option, you can configure the source IP addresses to failover to the target. If your target is on the same subnet as the source (typical of a LAN environment), you should select this option. Do not select this option if you are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

  

  If you are applying the source network configuration to the target in a WAN environement, do not failover your IP addresses 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 failover. If DNS updates are made, the DNS records will be locked during failover. Be sure and review the *Core Double-Take requirements* on page 23 for the requirements for updating DNS.

  

  DNS updates are not available for Server Core servers or NAT configurations.

  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.

If you select **Retain your target network configuration** but do not enable **Update DNS server**, you will need to specify failover scripts that update your DNS server during failover, or you can update the DNS server manually after failover. This would also apply to non--Microsoft Active Directory Integrated DNS servers. You will want to keep your target network configuration but do not update DNS. In this case, you will need to specify failover scripts that update your DNS server during failover, or you can update the DNS server manually after failover.

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

Reverse Protection

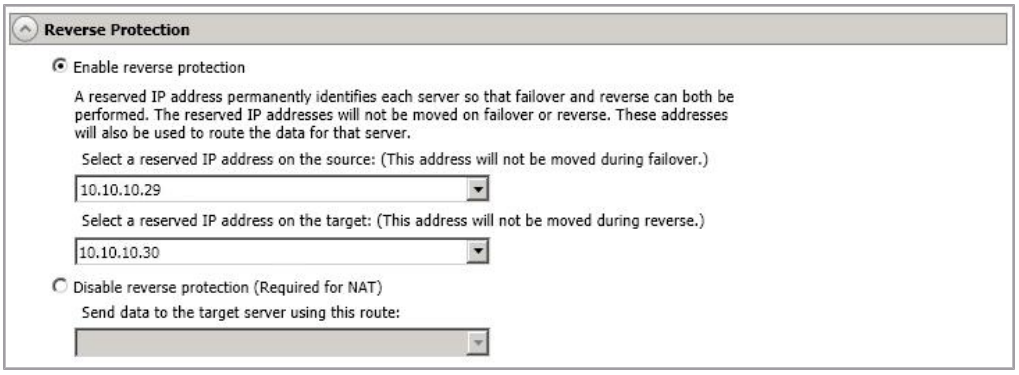

- **Enable reverse protection**—After failover, your target server is lost. Reverse protection allows you to store a copy of the target's system state on the source server, so that the target server will not be lost. The reverse process will bring the target identity back on the source hardware and establish protection. After the reverse, the source (running on the original target hardware) will be protected to the target (running on the original source hardware).

  In a LAN environment, you may want to consider having two IP addresses on each server. This will allow you to monitor and failover one (or more) IP addresses, while still leaving an IP address that does not get failed over. This IP address that is not failed over is called a reserved IP address and can be used for the reverse process. The reserved IP address remains with the server hardware. Ideally, the reserved IP address should not be used for production communications. The reserved IP address can be on the same or a different subnet from your production IP addresses, however if the subnet is different, it should be on a different network adapter. The reserved IP addresses will also be used to route Double-Take data.

  You do not have to have a second IP address on each server. It is acceptable to use the production IP address for reverse protection. In this case, Double-Take will block the DNS record for that address while it is failed over.

  - **Select a reserved IP address on the source**—Specify an IP address on the source which will be used to permanently identify the source server. The IP address you specify will not be failed over to the target in the event of a failure. This allows you to reverse protection back to the source after a failover.
  - **Select a reserved IP address on the target**—Specify an IP address on the target which will be used to permanently identify the target server. The IP address you specify will not be lost during failover. This allows you to reverse protection back to the source after a failover.

  

  When reverse protection is enabled, an image of the target's system state is mirrored to the source server. This mirror may cause a performance impact on your source server. This impact is only temporary, and system performance will return to normal when the reverse protection mirror is complete. To maintain system performance on the source, the target's system state is not continuously replicated to the source. You can manually update the image of the target's system state by

viewing the job details and clicking **Update** under **Target Server Image**. See *Viewing full server job details* on page 273.

- **Disable reverse protection**—If you do not use reverse protection, after a failover, your target server will be lost. In order to continue protecting your data, you will have to manually rebuild your original source and restart protection, which can be a long and complicated process. Also, if you disable reverse, you will lose the activated target license after failover. See *Deactivating licenses* on page 58 for how you can save and deactivate the target license. This is your only supported option if you are using a NAT environment.
  - **Send data to the target server using this route**—Specify an IP address on the target to route Double-Take data. 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 address of the NAT router) if you are using a NAT environment.

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.

Network Adapter Options

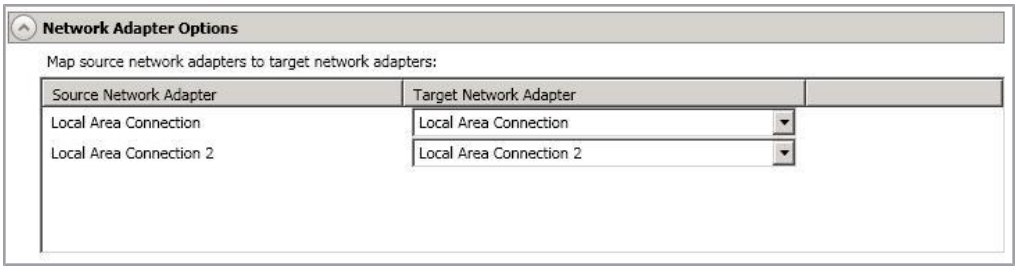

For **Map source network adapters to target network adapters**, specify how you want the IP addresses associated with each NIC on the source to be mapped to a NIC on the target. Do not mix public and private networks. Also, if you have enabled reverse protection, make sure that your NICs with your reserved IP addresses are mapped to each other.

Mirror, Verify & Orphaned Files

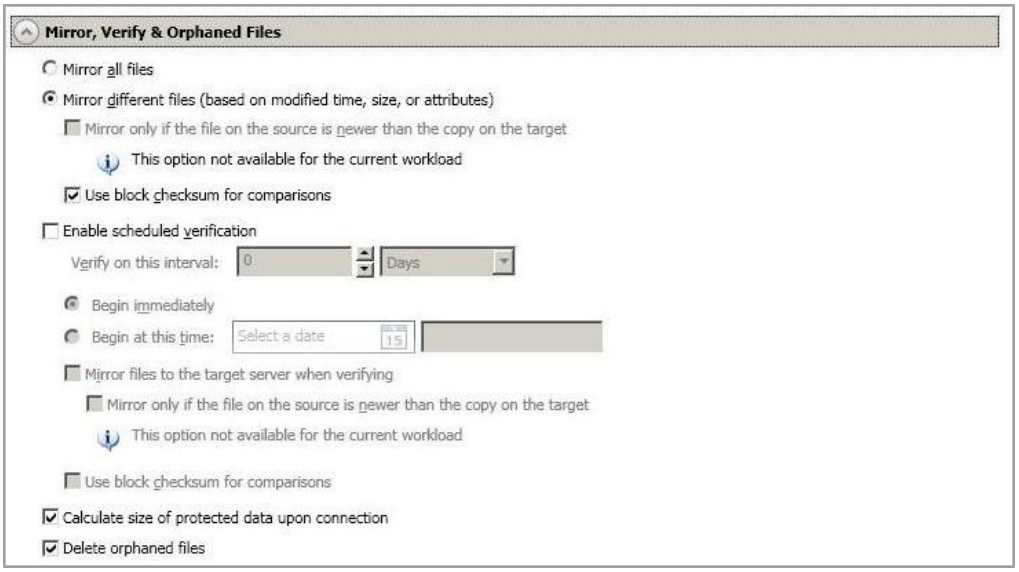

- **Mirror all files**—All protected files will be mirrored from the source to the target.
- **Mirror different files**—Only those protected files that are different based on date and time, size, or attributes will be mirrored from the source to the target.
  - **Mirror only if the file on the source is newer than the copy on the target**—This option is not available for full server jobs.
  - **Use block checksum for comparisons**—For those files flagged as different, the mirroring process can perform a block checksum comparison and send only those blocks that are different. Make sure you always enable this option for full server jobs.

File Differences Mirror Options Compared

The following table will help you understand how the various difference mirror options work together, including when you are using the block checksum option configured through the *Source server properties* on page 92.

An X in the table indicates that option is enabled. An X enclosed in parentheses (X) indicates that the option can be on or off without impacting the action performed during the mirror.

Not all job types have the source newer option available.

Source Server Properties	Job Properties			Action Performed
Block Checksum Option	File Differences Option	Source Newer Option	Block Checksum Option	
(X)	X			Any file that is different on the source and target based on the date, time, size, and/or attribute is transmitted to the target. The mirror sends the entire file.
(X)	X	X		Any file that is newer on the source than on the target based on date and/or time is transmitted to the target. The mirror sends the entire file.
	X		X	Any file that is different on the source and target based on date, time, size, and/or attributed is flagged as different. The mirror then performs a checksum comparison on the flagged files and only sends those blocks that are different.
X	X		X	The mirror performs a checksum comparison on all files and only sends those blocks that are different.
(X)	X	X	X	Any file that is newer on the source than on the target based on date and/or time is flagged as different. The mirror then performs a checksum comparison on the flagged files and only sends those blocks that are different.

- **Enable scheduled verification**—Verification is the process of confirming that the source replica data on the target is identical to the original 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, can automatically initiate a remirror, if configured. The remirror ensures data integrity between the source and target. When this option is enabled, Double-Take will verify the source replica data on the target and generate a verification log.

  

  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.

  - **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 at the specified date and time.
  - **Mirror files to the target server when verifying**—When this option is enabled, in addition to verifying the data and generating a log, Double-Take will also mirror to the target any protected files that are different on the source.
    - **Mirror only if the file on the source is newer than the copy on the target**—This option is not available for full server jobs.
    - **Use block checksum for comparisons**—For those files flagged as different, the mirroring process can perform a block checksum comparison and send only those blocks that are different.
- **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 the calculation is enabled, it is completed before the job starts mirroring, which can take a significant amount of time depending on the number of files and system performance. If your job contains a large number of files, for example, 250,000 or more, you may want to disable the calculation so that data will start being mirrored sooner. 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 system state protection processes hard links. If you disable the calculation, the hard link processing will not occur and you may have problems after failover, especially if your source is Windows 2008 or 2012.

- **Delete orphaned files**—An orphaned file is a file that exists in the replica data on the

target, but does not exist in the protected data on the source. This option specifies if orphaned files should be deleted on the target during a mirror, verification, or restoration.

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

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 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.

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 95 for more information.

During a mirror, orphaned file processing success messages will be logged to a separate orphaned file log. 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 100 for details on the location of that folder. The orphaned log file is overwritten during each orphaned file processing during a mirror, and the log file will be a maximum of 50 MB.

Staging Folder Options

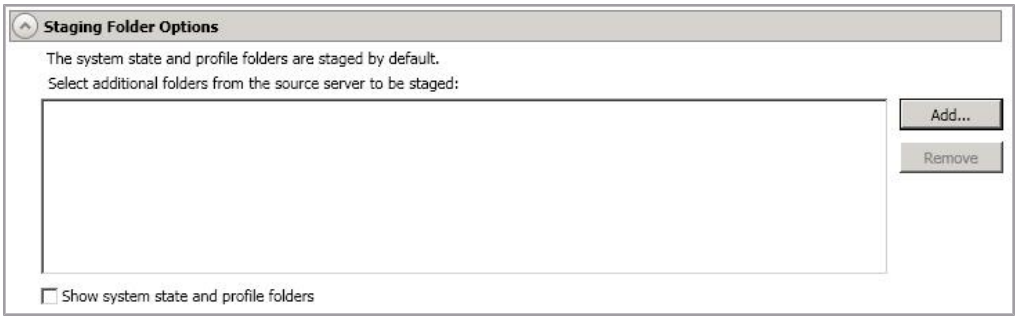

- **Select additional folders from the source that need to be staged**—Applications running on the target that cannot be stopped will cause retry operations because Double-Take will be unable to write to open application files. In this case, you will want to mirror those application files to a staging location instead of their actual location. Generally, this will only apply to applications that are not installed in the Windows Program Files directory. In this case, click **Add** and specify the folder that you want staged. Any staged folders will be applied to their actual installation location during failover.

  If IIS is being used as a software application on your source but as a hardware platform manager by your hardware vendor on your target, you need to add the INetPub directory to the **Staged Folders Options** list. If IIS is being used as a hardware platform manager by your hardware vendor on both the source and target, you need to go to the **Choose Data** page and remove the INetPub directory from replication under the **Replication Rules** heading.

- **Show system state and profile folders**—This option displays the list of essential system state and profile folders that will be staged automatically. These essential items are displayed in a lighter color than folders you have manually added, and they cannot be removed from the list.

Target Services

- **Services to leave running on the target server during protection**—Double-Take controls which services are running and stopped on the target during protection. You can specify which services you want to keep running by clicking **Add** and selecting a service from the list. If you want to remove a service from the list, highlight it and click **Remove**.

  Services are stopped on the target to protect against retry operations. Do not leave services running unless absolutely necessary.

- **Show essential services**—This option displays the list of essential services that will remain running on the target. The essential services are displayed in a lighter color than services you have manually added. The essential services cannot be removed from the list.

Snapshots

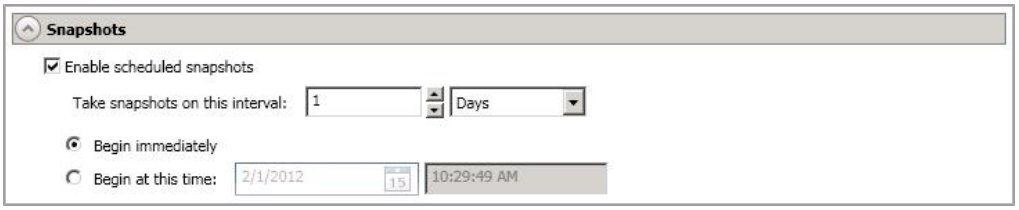

A snapshot is an image of the source replica data on the target taken at a single point in time. You can failover to a snapshot. However, you cannot access the snapshot to recover specific files or folders.

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 starting at a later date and time. Specify the date and time parameters to indicate when you want to start.

See *Managing snapshots* on page 161 for details on taking manual snapshots and deleting snapshots.

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

Compression

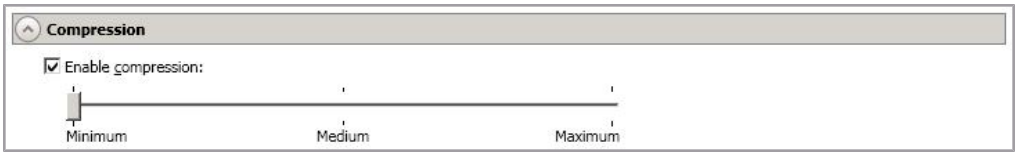

To help reduce the amount of bandwidth needed to transmit Double-Take data, compression allows you to compress data prior to transmitting it across the network. In a WAN environment this provides optimal use of your network resources. If compression is enabled, the data is compressed before it is transmitted 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

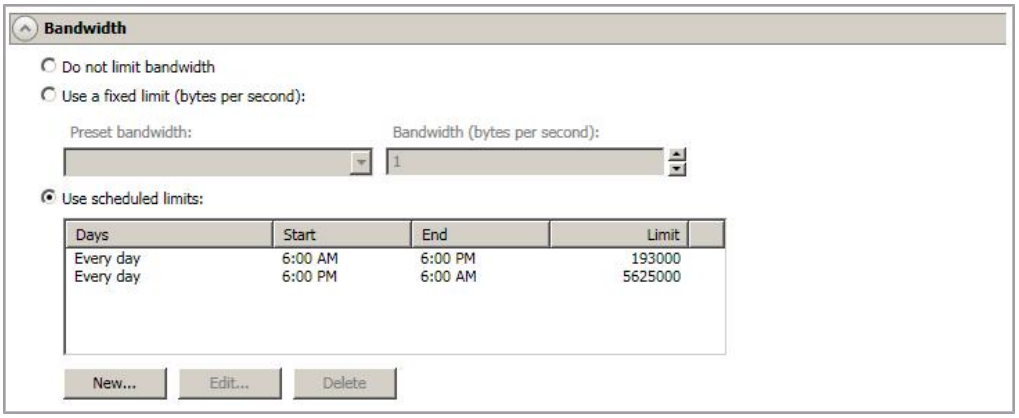

Bandwidth limitations are available to restrict the amount of network bandwidth used for Double-Take data transmissions. When a bandwidth limit is specified, Double-Take never exceeds that allotted amount. The bandwidth not in use by Double-Take is available for all other network traffic.

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

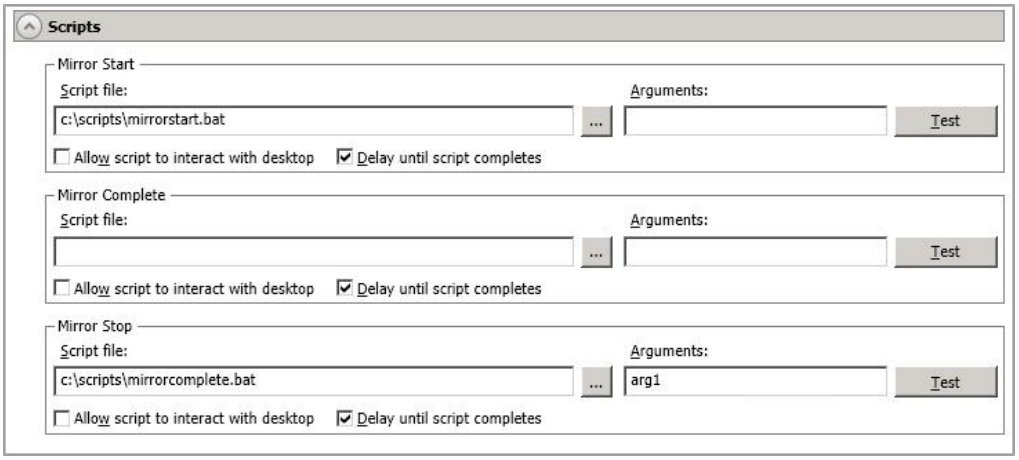

You can customize mirroring by running scripts on the target at pre-defined points in the mirroring process. Scripts may contain any valid Windows command, executable, or batch file. The scripts are processed using the same account running the Double-Take service, unless you have identified a specific account through the server's properties. See *Script credentials* on page 99. 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.
- **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.

Mirror scripts are dependent on the target and the **Target Path Mappings** specified under the **Network Route & Folder Selection** section. 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).

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

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

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

Jobs in a NAT environment may take longer to start.

Managing and controlling full server jobs

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

The jobs displayed in the right pane depend on the server group folder selected in the left pane. Every job for each server in your console session is displayed when the **Jobs on All Servers** group is selected. If you have created and populated server groups (see *Managing servers* on page 65), then only the jobs associated with the server or target servers in that server group will be displayed in the right pane.

- See *Overview job information displayed in the top pane* on page 264
- See *Detailed job information displayed in the bottom pane* on page 266
- See *Job controls* on page 268

Overview job information displayed in the top pane

The top pane displays high-level overview information about your jobs.

Column 1 (Blank)

The first blank column indicates the state of the job.

 The job is in a healthy state.

 The job is in a warning state. This icon is also displayed on any server groups that you have created that contain a job in a warning state.

 The job is in an error state. This icon is also displayed on any server groups that you have created that contain a job in an error state.

 The job is in an unknown state.

Job

The name of the job

Source Server

The name of the source. This could be the name or IP address of your source. In parenthesis will be the reserved IP address that you assigned to this server.

Target Server

The name of the target. This could be the name or IP address of your target. In parenthesis will be the reserved IP address that you assigned to this server.

Job Type

Each job type has a unique job type name. This job is a Full Server Failover job. For a complete list of all job type names, press F1 to view the Double-Take Console online help.

Activity

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

Mirror Status

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

Replication Status

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

Transmit Mode

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

Detailed job information displayed in the bottom pane

The details displayed in the bottom pane of the **Manage Jobs** page provide additional information for the job highlighted in the top pane. If you select multiple jobs, the details for the first selected job will be displayed.

Name

The name of the job

Target data state

- **OK**—The data on the target is in a good state.
- **Mirroring**—The target is in the middle of a mirror process. The data will not be in a good state until the mirror is complete.
- **Mirror Required**—The data on the target is not in a good state because a remirror is required. This may be caused by an incomplete or stopped mirror or an operation may have been dropped on the target.
- **Restore Required**—The data on the source and target do not match because of a failover condition. Restore the data from the target back to the source. If you want to discard the changes on the target, you can remirror to resynchronize the source and target.
- **Snapshot Reverted**—The data on the source and target do not match because a snapshot has been applied on the target. Restore the data from the target back to the source. If you want to discard the changes on the target, you can remirror to resynchronize the source and target.
- **Busy**—The source is low on memory causing a delay in getting the state of the data on the target.
- **Not Loaded**—Double-Take target functionality is not loaded on the target server. This may be caused by an activation code error.
- **Unknown**—The console cannot determine the status.

Mirror remaining

The total number of mirror bytes that are remaining to be sent from the source to the target

Mirror skipped

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

Replication queue

The total number of replication bytes in the source queue

Disk queue

The amount of disk space being used to 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**.

Connected since

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

Recent activity

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

Additional information

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

Job controls

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

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

Create a New Job 

This button leaves the **Manage Jobs** page and opens the **Get Started** page.

View Job Details 

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

Delete 

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 77. You will remain on the **Manage Jobs** page after updating the server credentials. If your servers use the same credentials, make sure you also update the credentials on the **Manage Servers** page so that the Double-Take Console can authenticate to the servers in the console session. See *Managing servers* on page 65.

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. Failover monitoring will continue 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. Failover monitoring will continue while the job is stopped. 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 161 for more information.

Snapshots are not applicable to clustered environments.

Failover or Cutover 

Starts the failover process. See *Failing over full server jobs* on page 282 for the process and details of failing over a full server job.

Failback 

Starts the failback process. Failback does not apply to full server jobs.

Restore 

Starts the restoration process. Restore does not apply to full server jobs.

Reverse 

Reverses protection. The original source hardware will be reversed to the target identity and the job will start mirroring in the reverse direction with the job name and log file names changing accordingly. After the mirror is complete, the job will continue running in the opposite direction. See *Reversing full server jobs* on page 285 for the process and details of reversing a full server job.

Undo Failover 

Cancels a test failover by undoing it. Undo failover does not apply to full server jobs.

View Job Log 

Opens the job log. On the right-click menu, this option is called **View Logs**, and you have the option of opening the job log, source server log, or target server log. See *Viewing the log files through the Double-Take Console* on page 678 for details on all three of these logs.

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 all files**—All protected files will be mirrored from the source to the target.
  - **Mirror different files**—Only those protected files that are different based on date and time, size, or attributes will be mirrored from the source to the target.
    - **Mirror only if the file on the source is newer than the copy on the target**—This option is available for full server jobs, but ideally it should not be used.
    - **Use block checksum for comparisons**—For those files flagged as different, the mirroring process can perform a block checksum comparison and send only those blocks that are different. Make sure you always enable this option for full server jobs.
  - **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 the calculation is enabled, it is completed before the job starts mirroring, which can take a significant amount of time depending on the number of files and system performance. If your job contains a large number of files, for example, 250,000 or more, you may want to disable the calculation so that data will start being mirrored sooner. 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 system state protection processes hard links. If you disable the calculation, the hard link processing will not occur and you may have problems after failover, especially if your source is Windows 2008.

- **Verify**—Even if you have scheduled the verification process, you can run it manually any time a mirror is not in progress.
  - **Create verification report only**—This option verifies the data and generates a verification log, but it does not remirror any files that are different on the source and target. See *Verification log* on page 102 for details on the log file.
  - **Mirror files to the target server automatically**—When this option is enabled, in addition to verifying the data and generating a log, Double-Take will also mirror to the target any protected files that are different on the source.
    - **Mirror only if the file on the source is newer than the copy on the target**—This option is available for full server jobs, but ideally it should not be used.
    - **Use block checksum for comparisons**—For those files flagged as different, the mirroring process can perform a block checksum comparison and send only those blocks that are different.
- **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 paused your target, failover will not start if configured for automatic failover, and it cannot be initiated if configured for manual intervention. You must resume the target before failover will automatically start or before you can manually start it.

  If you have multiple jobs to the same target, all jobs from the same source will be paused and resumed.
- **Update Shares**—Shares are not applicable because they are automatically included with the system state that is being protected with the entire server.

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 65.

Type a server name

Displays only jobs that contain the text you entered. If you have created and populated server groups, then only jobs that contain the text you entered associated with the server or target servers in that server group will be displayed. See *Managing servers* on page 65.

Overflow Chevron 

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

Viewing full server job details

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

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

Job name

The name of the job

Job type

Each job type has a unique job type name. This job is a Full Server Failover job. For a complete list of all job type names, press F1 to view the Double-Take Console online help.

Health

 The job is in a healthy state.

 The job is in a warning state.

 The job is in an error state.

 The job is in an unknown state.

Activity

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

Connection ID

The incremental counter used to number connections. The number is incremented when a connection is created. It is also incremented by internal actions, such as an auto-disconnect and auto-reconnect. The lowest available number (as connections are created, stopped, deleted, and so on) will always be used. The counter is reset to one each time the Double-Take service is restarted.

Transmit mode

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

Target data state

- **OK**—The data on the target is in a good state.
- **Mirroring**—The target is in the middle of a mirror process. The data will not be in a good state until the mirror is complete.
- **Mirror Required**—The data on the target is not in a good state because a remirror is required. This may be caused by an incomplete or stopped mirror or an operation may have been dropped on the target.
- **Restore Required**—The data on the source and target do not match because of a failover condition. Restore the data from the target back to the source. If you want to discard the changes on the target, you can remirror to resynchronize the source and target.
- **Snapshot Reverted**—The data on the source and target do not match because a snapshot has been applied on the target. Restore the data from the target back to the source. If you want to discard the changes on the target, you can remirror to resynchronize the source and target.
- **Busy**—The source is low on memory causing a delay in getting the state of the data on the target.
- **Not Loaded**—Double-Take target functionality is not loaded on the target server. This may be caused by an activation code error.
- **Unknown**—The console cannot determine the status.

Target route

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

Compression

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

Encryption

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

Bandwidth limit

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

Connected since

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

Additional information

Depending on the current state of your job, you may see additional information displayed to keep you informed about the progress and status of your job. If there is no

additional information, you will see (None) displayed.

Mirror status

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

Mirror percent complete

The percentage of the mirror that has been completed

Mirror remaining

The total number of mirror bytes that are remaining to be sent from the source to the target

Mirror skipped

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

Replication status

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

Replication queue

The total number of replication bytes in the source queue

Disk queue

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

Bytes sent

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

Bytes sent compressed

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

Target Server Image

When a full server job is created with reverse protection enabled, an image of the target's system state is stored on the source server. This image allows you to reverse your source and target after a failover. To improve performance, the target's system state is not continuously replicated to the source. You should manually update the image of the target's system state by clicking **Update** if there is a change on the target. For example, if the credentials on the target server are updated, you should update the target server image that is on the source. This reverse protection mirror may cause a performance impact on your source server. This impact is only temporary, and system performance will return to normal when the reverse protection mirror is complete.

Validating a full server job

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

1. From the **Manage Jobs** page, highlight the job and click **View Job Details** in the toolbar.
2. In the **Tasks** area on the right on the **View Job Details** page, click **Validate job properties**.
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 enable protection. Depending on the error, you may be able to click **Fix** or **Fix All** and let Double-Take correct the problem for you. For those errors that Double-Take cannot correct automatically, you will need to modify the source or target to correct the error, or you can select a different target. You must revalidate the selected servers, by clicking **Recheck**, until the validation check passes without errors.

   Validation checks for an existing job are logged to the job log on the target server. See *Log files* on page 677 for details on the various log files.
4. Once your servers have passed validation, click **Close**.

Editing a full server job

Use these instructions to edit a full server job.

1. From the **Manage Jobs** page, highlight the job and click **View Job Details** in the toolbar.
2. In the **Tasks** area on the right on the **View Job Details** page, click **Edit job properties**. (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 for your full server job as when you created the job, but you will not be able to edit all of them. If desired, edit those options that are configurable for an existing job. See *Creating a full server job* on page 238 for details on each job option.

   

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

4. If you want to modify the workload items or replication rules for the job, click **Edit workload or replication rules**. Modify the **Workload item** you are protecting, if desired. Additionally, you can modify the specific **Replication Rules** for your job.

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

   You can also enter wildcard rules, however you should do so carefully. Rules are applied to files that are closest in the directory tree to them. If you have rules that include multiple folders, an exclusion rule with a wild card will need to be added for each folder that it needs applied to. For example, if you want to exclude all .log files from D:\ and your rules include D:\, D:\Dir1 , and D:\Dir2, you would need to add the exclusion rule for the root and each subfolder rule. So you will need to add exclude rules for D:*.log , D:\Dir1*.log, and D:\Dir2*.log.

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

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

   Errors are designated by a white X inside a red circle. Warnings are designated by a black exclamation point (!) inside a yellow triangle. A successful validation is designated by a white checkmark inside a green circle. You can sort the list by the icon to see errors, warnings, or successful validations together. Click on any of the validation items to see details. You must

correct any errors before you can enable protection. Depending on the error, you may be able to click **Fix** or **Fix All** and let Double-Take correct the problem for you. For those errors that Double-Take cannot correct automatically, you will need to modify the source or target to correct the error, or you can select a different target. You must revalidate the selected servers, by clicking **Recheck**, until the validation check passes without errors.

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.

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

Viewing a full server job log

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

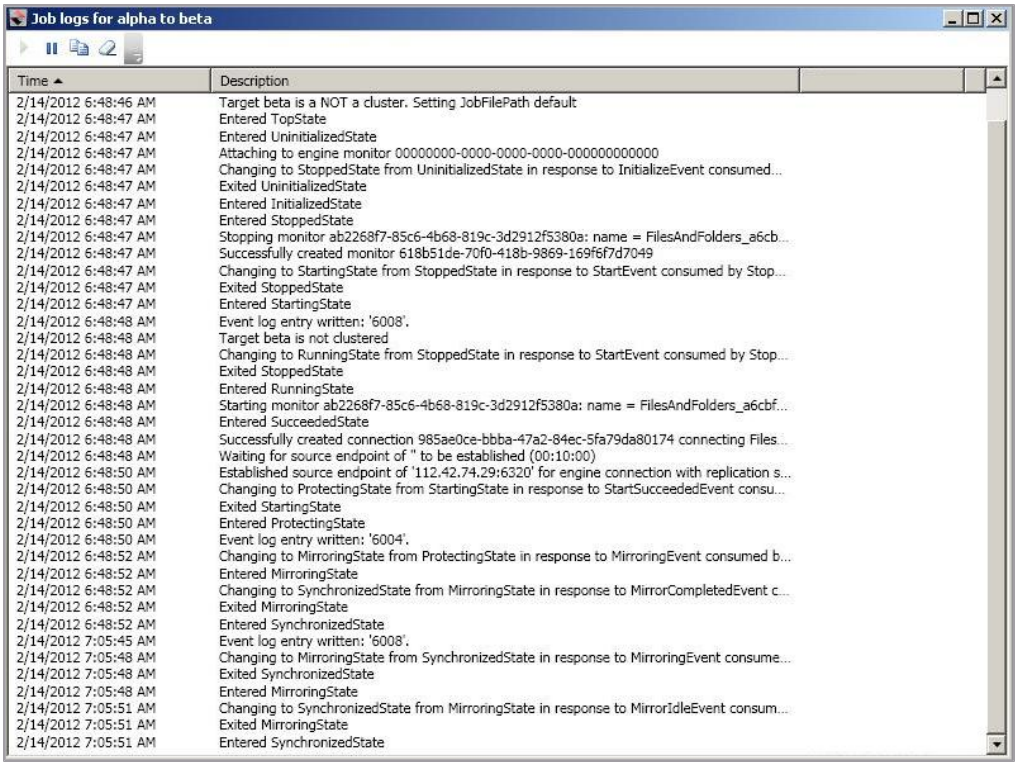

The following table identifies the controls and the table columns in the **Job logs** window.

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.

Failing over full server jobs

When a failover condition has been met, failover will be triggered automatically if you disabled the wait for user option during your failover configuration. If the wait for user before failover option is enabled, you will be notified in the console when a failover condition has been met. At that time, you will need to trigger it manually from the console when you are ready.

If you have paused your target, failover will not start if configured for automatic failover, and it cannot be initiated if configured for manual intervention. You must resume the target before failover will automatically start or before you can manually start it.

1. On the **Manage Jobs** page, highlight the job that you want to failover and click **Failover or Cutover** in the toolbar.
2. Select the type of failover to perform.
   - **Failover to live data**—Select this option to initiate a full, live failover using the current data on the target. The source is automatically shut down if it is still running. Then the target will stand in for the source by rebooting and applying the source identity, including its system state, on the target. After the reboot, the target becomes the source, and the target no longer exists.
   - **Perform test failover**—This option should only be used if your target is a virtual server. It is like live failover, except the source is not shutdown. Therefore you should isolate the virtual server from the network before beginning the test using the following procedure.
     a. Stop the job.
     b. Take a snapshot of the target virtual server.
     c. Attach the target virtual server to a null virtual switch or one that does not have access to your network infrastructure.
     d. Perform the test failover and complete any testing on the virtual server.
     e. After your testing is complete, revert to the snapshot of the target virtual server from before the test started.
     f. Reconnect the target virtual server to the proper virtual switch.
     g. Restart the job.

     If your target is a physical server, contact technical support if you want to test failover, because you will have to rebuild your target system volume after the test.
   - **Failover to a snapshot**—Select this option to initiate a full, live failover without using the current data on the target. Instead, select a snapshot and the data on the target will be reverted to that snapshot. This option will not be available if there are no snapshots on the target or if the target does not support snapshots. This option is also not applicable to clustered environments. To help you understand what snapshots are available, the **Type** indicates the kind of snapshot.
     - **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. Select how you want to handle the data in the target queue. You may want to check the amount of data in queue on the target by reviewing the *Statistics* on page 688 or *Performance Monitor* on page 790.
   - **Apply data in target queues before failover or cutover**—All of the data in the target queue will be applied before failover begins. The advantage to this option is that all of the data that the target has received will be applied before failover begins. 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 failover or cutover immediately**—All of the data in the target queue will be discarded and failover will begin immediately. The advantage to this option is that failover will occur immediately. The disadvantage is that any data in the target queue will be lost.
   - **Revert to last good snapshot if target data state is bad**—If the target data is in a bad state, Double-Take will automatically revert to the last good Double-Take snapshot before failover begins. If the target data is in a good state, Double-Take will not revert the target data. Instead, Double-Take will apply the data in the target queue and then failover. The advantage to this option is that good data on the target is guaranteed to be used. The disadvantage is that if the target data state is bad, you will lose any data between the last good snapshot and the failure.
4. When you are ready to begin failover, click **Failover**.

IPv6 addresses on the source will be set to DHCP on the target after failover. Update them to static addresses manually, if needed.

You may experience issues following a failover if an application or server uses hard-linked files. For example, Windows 2008 or 2012 Server Roles added after the job has been established will not function after failover because the hard links related to the server role were not replicated. After updating server roles, a remirror should be performed.

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 failed over because it appears to be a hardware device.

If your NICs were configured for network load balancing (NLB), you will have to reconfigure that after failover.

Because the Windows product activation is dependent on hardware, you may need to reactivate your Windows registration after failover. In most cases when you are using Windows 2003, you can follow the on-screen prompts to complete the reactivation. However, when you are using Windows 2008 or 2012, the reactivation depends on several factors including service pack level, Windows edition, and your licensing type. If a Windows 2008 or 2012 target comes online after failover 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.

5. If you performed a test failover, revert your snapshot on the target, and then restart your job by clicking **Start**.

Reversing full server jobs

After a full server failover, the source is running on your original target hardware and your target no longer exists. That means the source and target hardware now share the same identity, which is the source identity.

If you did not enable reverse protection, your source is a domain controller, or if you have to rebuild your source, you will have to reverse your protection manually. See *Reversing full server jobs manually* on page 287.

---

1. Fix the issue that caused your original source server to fail.
2. Connect the original source server to the network.
3. Make sure the production NIC on your original source is online. If the NIC is disabled or unplugged, you will experience problems with the reverse. Make sure you continue to access the servers through the reserved IP addresses, but you can disregard any IP address conflicts for the primary NIC. Since the new source (running on the original target hardware) already has the source's address assigned to it, Windows will automatically assign a different address to the original source.
4. On the **Manage Jobs** page, highlight the job that you want to reverse. If the job is not listed, you may need to add your servers to your console again. Use the reserved IP addresses and local credentials.
5. Highlight the job you want to reverse and click **Reverse** in the toolbar. During the reverse process, you will see various states for the job. During the **Restoring** process, the target identity is being established on the original source hardware. During the **Synchronizing** process , protection is being established from the source (on the original target hardware) to the target (on the original source hardware). The reverse protection is also established in the opposite direction. When the reverse process is complete, the target (on the original source hardware) will reboot. At this point, your source is still running on your original target hardware with the source name, but the original source hardware now has the target identity.
6. To go back to your original hardware, highlight the job and click **Failover**. The source identity will now be applied to the target (on the original source hardware), and the target identify will again be gone. Both servers will have the source identity.
7. To bring back the target identify, highlight the job and click **Reverse**. The same process as above will be repeated, but on the opposite servers. When the reverse is complete, you will be back to your original identities on the original hardware.

IPv6 addresses on the source will be set to DHCP on the target after failover. Update them to static addresses manually, if needed.

You may experience issues following a failover if an application or server uses hard-linked files. For example, Windows 2008 or 2012 Server Roles added after the job has been established will

not function after failover because the hard links related to the server role were not replicated. After updating server roles, a remirror should be performed.

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 failed over because it appears to be a hardware device.

If your NICs were configured for network load balancing (NLB), you will have to reconfigure that after failover.

Because the Windows product activation is dependent on hardware, you may need to reactivate your Windows registration after failover. In most cases when you are using Windows 2003, you can follow the on-screen prompts to complete the reactivation. However, when you are using Windows 2008 or 2012, the reactivation depends on several factors including service pack level, Windows edition, and your licensing type. If a Windows 2008 or 2012 target comes online after failover 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.

Reversing full server jobs manually

If you did not enable reverse protection, your source is a domain controller, or if you have to rebuild your source, you have two options after a failover. You can continue running from the failed over server indefinitely. This server is your source (running on the original target hardware) and you can protect it to a new target. Your other option is to go back to the original hardware. Without reverse protection, you have to complete this process manually, which can be difficult.

Preparation of your original source hardware or a new server is key to this manual process. The type of preparation required will depend on the role of the original source server, the applications that were used on the original server, whether the original source was a physical or virtual server, and the failure or event that occurred.

- **Server role**—If your original source was a domain controller, a Cluster Service server, or a Certificate Service server, you will have to reinstall Windows. The utility required to reuse a server cannot be used on these types of servers. Start with *1A. Preparing a new server by reinstalling Windows* on page 288 and then continue with the remaining instructions.
- **Applications**—If your original source was running a name-specific application, like Exchange, you should reinstall Windows. Start with *1A. Preparing a new server by reinstalling Windows* on page 288 and then continue with the remaining instructions.
- **Physical servers**—If your original source was a physical server, your preparation method will depend on if you experienced a catastrophic or non-catastrophic failure.
  - **Catastrophic failure**—If your original hardware is unusable or the failure will require you to reinstall Windows, start with *1A. Preparing a new server by reinstalling Windows* on page 288 and then continue with the remaining instructions.
  - **Non-catastrophic failure**—If the failure did not damage the server or the operating system and you want to reuse the server, start with *1B. Reusing your original source hardware* on page 288 and then continue with the remaining instructions.
- **Virtual servers**—If your original source was a virtual server, you preparation method will depend on if you want to create a new virtual guest or reuse the existing one.
  - **New virtual guest**—If your original guest is unusable, start with *1A. Preparing a new server by reinstalling Windows* on page 288 and then continue with the remaining instructions.

    

    If possible, you can attach any virtual hard disks that survived the failure event to a new virtual guest. Reusing any undamaged disks will decrease the time required to restore data because you can use a difference mirror.

    As an alternative to manually creating a new virtual guest, you can let Double-Take automatically provision (create) the new virtual guest for you. If you choose this option, you will need to use the instructions *Creating a full server to ESX job* on page 445 or *Creating a full server to Hyper-V job* on page 397 instead of the instructions in this section.

  - **Reusing virtual guest**—If the failure did not damage the virtual guest and you want to reuse it, start with *1B. Reusing your original source hardware* on page 288 and then continue with the remaining instructions.

1A. Preparing a new server by reinstalling Windows

1. Install or reinstall Windows on your physical or virtual server using unique, temporary server information. See your Windows documentation for details on installing the operating system.
2. After the operating system installation is complete, install Double-Take using the activation code from your original target. See *Installing using the installation wizard* on page 37.
3. After Double-Take is installed and activated, continue with *2. Mirroring and replicating from the source to the original source hardware or new server and failing over* on page 289.

1B. Reusing your original source hardware

1. Disconnect the original source hardware from the network. For a physical server, you may want to disconnect the network cable. For a virtual server, remove it from the network using your virtual console. You must make sure the original source is completely disconnected before proceeding.
2. After the original source hardware is disconnected from the network, remove the target server identity from Active Directory. You should remove the target's original identity, not the identity of the source which the original target hardware now holds.
3. Keeping the original source hardware disconnected from the network, reboot it and login as the local administrator.
4. Stop all application services on the original source hardware and set them to manual.
5. If you failed over the source IP address, create a new unique IP addresses. See your Windows documentation for details on modifying IP addresses.
6. Modify the original source hardware identity by placing the server into a workgroup. Make sure you reboot when prompted, continuing to keep the server disconnected from the network. See your Windows documentation for details on placing a server into a workgroup.
7. After the reboot, login as the local administrator.
8. Using the Double-Take Console, remove and reinsert the original source server into the server list on the **Manage Servers** page.
9. Double-click on the server in the server list to view the server details page, and then click on the **Edit server properties** link.
10. Under the **Licensing** section, enter the activation code from the original target server and click **Add**. If the original source server activation code is listed, remove it from the **Current activation codes** list. Click **OK** to return to the **Manage Servers** page.
11. Run the Microsoft Sysprep utility to modify SIDs (security identifiers) and the server name. If desired, you can use the original target server name when the utility prompts for a server name. See the Microsoft web site for details on the Sysprep utility.

If the Sysprep utility does not force you to choose a new computer name, you will need to complete the following additional steps.

1. Finish the Sysprep process.
2. Reboot the server and login as the local administrator.
3. Rename the computer manually and reboot when prompted.

The server must be given a new name either via Sysprep or manually after Sysprep has completed before you can proceed.

12. Connect the server to the network and continue with *2. Mirroring and replicating from the source to the original source hardware or new server and failing over* on page 289.

2. Mirroring and replicating from the source to the original source hardware or new server and failing over

1. Using the Double-Take Console, delete the original job from the **Manage Jobs** page.
2. Establish full server protection from your source to the original source hardware or new server. See *Creating a full server job* on page 238. In the console, specify your source (running on the original target hardware) on the **Choose Source Server** page, and specify your original source hardware or new server that you built or modified in step 1A or 1B above on the **Select Target Server** page. Select the same data for protection and use the options that you used when protecting the source initially, although you can select different settings for snapshots, compression, and so on.
3. Once you have established full server protection, data will be mirrored from the source (on the original target hardware) to the target (on the original source hardware or your new server). Replication will keep the target up-to-date with the changes end-users are continuing to make on the source. Monitor the progress of the job. See *Managing and controlling full server jobs* on page 264.
4. Once the mirror is complete, determine when you want to perform failover. This will require downtime, typically between 15 and 30 minutes depending on LAN or WAN configurations and server processing capabilities.
5. Using the Double-Take Console, perform failover using live data or a snapshot, as desired. See *Failing over full server jobs* on page 282.
6. Monitor the progress. After the target reboots, the target will no longer exist, since it will become the source.

After the reboot, users and other servers can resume normal operations after DNS/IP updates have been propagated to them.

If desired, you can re-establish protection again for this source so that you are prepared for the next emergency.

If you want to reuse the same target hardware, you will have to remove the source identity components from that server. You can use either method *1A. Preparing a new server by reinstalling Windows* on page 288 or *1B. Reusing your original source hardware* on page 288 to prepare a new target.

Chapter 9 Exchange protection

Create an Exchange job when you have Microsoft Exchange and want application-level protection.

- See *Exchange requirements* on page 291—Exchange protection includes specific requirements for this type of protection.
- See *Creating an Exchange job* on page 295—This section includes step-by-step instructions for creating a Exchange job.
- See *Managing and controlling Exchange jobs* on page 321—You can view status information about your Exchange jobs and learn how to control these jobs.
- See *Failing over Exchange jobs* on page 339—Use this section when a failover condition has been met or if you want to failover manually.
- *Restoring then failing back Exchange jobs* on page 341—Use this section when you are ready to restore and failback.

If your source is a domain controller, you should use one of the full server protection methods to protect the entire server because of complexities with authentication. See *Selecting a protection type* on page 165.

Exchange requirements

In addition to the *Core Double-Take requirements* on page 23, you must also meet the following requirements to protect Exchange.

- **Exchange versions**—Double-Take can protect Microsoft Exchange 2003, 2007 or Exchange 2010. Exchange 2010 must have Service Pack 1 or later.
- **Exchange and network configuration**—The following requirements and limitations apply to your Exchange server and network configuration.
  - All Microsoft best practices should be used for all versions of Exchange.
  - The version of Exchange on the source and target must be identical.
  - Windows 2012 is not supported.
  - Exchange 2010 must be running on a 64-bit server running Windows 2008 SP2 or later or Windows 2008 R2.
  - If you are protecting Exchange 2010 and you have a consolidated target server, you must have a send connector configured specifically with the target server before failover. Otherwise, you will be unable to send email to the Internet after failover.
  - If you are protecting Exchange 2010, arbitration mailboxes will not be failed over. These mailboxes can be rehomed manually using the Set-Mailbox -database PowerShell command.
  - Double-Take does not check the edition of Exchange 2007 (Enterprise or Standard). However, it is able to differentiate between service pack levels. If you have Exchange 2007 Enterprise on the source and Exchange 2007 Standard on the target, you will be limited to only failing over the number of databases or storage groups supported by Exchange 2007 Standard. See the Exchange Server 2007 Editions and Client Access Licenses information on the Microsoft website.
  - For Exchange 2007 and 2010, in a consolidated role environment only the mailbox role is protected. The Hub Transport and Client Access roles are not protected or failed over because they are already installed on the target.
  - For Exchange 2007 and 2010, replication and failover between a distributed role source configuration to a consolidated role target configuration is permitted as long as the source Mailbox Server role is installed on a standalone server or cluster with the other roles residing on different servers, and the target configuration is a standalone server with the Mailbox, Hub Transport, and Client Access roles installed. In these configurations, Double-Take will not replicate any data associated with the Hub Transport/Client Access data, however, the target Hub Transport/Client Access roles function properly when failing over the source Mailbox role, allowing necessary operations to resume. For Exchange 2010 with DAG (Database Availability Group), replication and failover from a distributed role source configuration must be to a non-DAG distributed role target configuration.
  - For Exchange 2010 with DAG, the following requirements and limitations also apply.
    - A DAG to standalone configuration is the only supported configuration. Multi-site DAGs (DAG to DAG configurations) are not supported.
    - All mailbox stores must be replicated to all other members of the DAG.
    - During failover, Double-Take will update SPNs, move user mailboxes, and perform Active Directory updates. DNS updates are not made during failover, but can be scripted manually if needed.

- Exchange 2003 on a domain controller is not a recommended configuration. However, if you must run Exchange 2003 on a domain controller, review Microsoft Knowledge Base articles 822179, 332097, 305065, 304403, and 875427.
- If you are protecting Exchange 2003 and it is running in mixed mode, the first installed Exchange virtual server contains the MTA (Message Transfer Agent) resource that is needed to communicate with versions prior to Exchange 2003. If you do not failover all Exchange virtual servers, then any user who is in a different mail store than the first one may not be able to route mail.
- The Exchange program files must be installed in the same location on the source and target.
- The drive letter(s) where Exchange stores its data on the source must be the same on the target.
- Single-label DNS domain names (those without a suffix such as .com, .corp, .net) are not supported.
- In environments where the FIPS security policy is enabled, you must use impersonation, which requires the following.
    - The user running the Double-Take Console must have all appropriate rights to update the domain (that is, only impersonation is supported).
    - You must manually verify DNS rights by running the DFO utility with the /test parameter.
- If you are protecting Exchange 2007 and will only be protecting a subset of the storage groups, the storage groups you are protecting cannot have a hyphen followed by a space in the storage group name. You will either need to change the name of the storage group to remove the hyphen followed by the space, or select all of your storage groups for protection.
- Microsoft Server Core is not supported.
- The source and target servers must be in the same root forest domain.
- In a parent/child domain, at least one domain controller in the child domain must be designated as a global catalog server.
- The target server cannot be a domain controller.
- Exchange and a domain controller cannot be on the same node of a cluster.
- The source and target servers must be part of the same Exchange Administrative Group.
- The Exchange configurations on the source and target servers must be identical for the following components: storage groups, location of storage groups (log and data files), log file prefixes, database locations (log and data files), Message Transfer Agent (MTA) location, and queue paths.
- Before you attempt to protect your Exchange application, you may want to complete the following tasks to verify that the environment is properly set up.
    - With both Exchange servers online, use **Active Directory Users and Computers** to move an existing user from the source to the target and then back to the original source.
    - Verify that you can create a new user on the target.
    - To verify connectivity, create an Outlook profile for the new user on a client machine and connect to the target.

- If /domainprep has not been run in an Exchange 2007 environment, users will not be failed over and SPNs will not be updated during failover due to access denied errors. To fix this issue, run setup with the /domainprep parameter in the environment.

- **Snapshots**—You can take and failover to Double-Take snapshots using an Exchange job, however snapshots are not supported if your source and/or target is a cluster. See *Core Double-Take requirements* on page 23 for the specific snapshot requirements.
- **Supported configurations**—The following table identifies the supported configurations for an Exchange job.

Configuration		Supported	Not Supported
Source to target configuration	One to one, active/standby	X	
	One to one, active/active		X
	Many to one		X
	One to many		X
	Chained		X
	Single server		X
Server configuration[1]	Standalone to standalone	X	
	Standalone to cluster		X
	Cluster to standalone	X	
	Cluster to cluster	X	
	Cluster Shared Volumes (CSV) guest level	X	
	Cluster Shared Volumes (CSV) host level		X
Upgrade configuration[2]	Upgrade 5.3 Double-Take Application Manager job to 7.0 Double-Take Console Exchange job		X
	Upgrade 6.0 Exchange job to 7.0 Exchange job	X	
Version 7.0 console[3]	Version 7.0 console can create job for 5.3 source and 5.3 target		X
	Version 7.0 console can create job for 6.0 source and 6.0 target	X	
	Version 7.0 console can create job for 7.0 source and 7.0 target	X	

1. When using a supported cluster configuration, the following limitations also apply.
   - If you are using Exchange 2007, only the mailbox role is protected.
   - Exchange and the domain controller cannot be on the same node in the cluster.

- Exchange must be installed in a unique group, not in the cluster group.

2. For an upgrade configuration that is not supported, you will have to delete the existing job before the upgrade and create a new job after the upgrade.
3. Newer job options available in the version 7.0 console will not be functional when creating jobs for servers running version 6.0.

Creating an Exchange job

Use these instructions to create an Exchange job.

1. Click **Get Started** from the toolbar.
2. Select **Double-Take Availability** and click **Next**.
3. Select **Protect files and folders, an application, or an entire Windows server** and click **Next**.
4. Choose your source server. This is the physical or virtual server running Exchange. If your source is a cluster, select the cluster name, not the Exchange virtual server name or virtual IP address. If you are protecting Exchange 2010 with DAG, select the DAG as your source server.

- **Current Servers**—This list contains the servers currently available in your console session. Servers that are not licensed for the workflow you have selected will be filtered out of the list. Select your source server from the list.
- **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 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.

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. If desired, the user must also

meet the following credentials requirements, however, you can specify a separate user later in the job creation process to handle the Exchange specific tasks.

- The account must be an Exchange Full Administrator at the organizational level, as delegated via the Exchange System Manager at the user level.
- The account must have rights to manage Exchange in order to query and modify the Exchange Active Directory objects.
- The account must be a member of the Domain Admins group. If your security policies do not allow use of this group, see *DNS* on page 816 and use the instructions under the Double-Take DFO utility to use a non-Domain Admins account.
- If Exchange is on a cluster, the account must be a member of the Cluster Administrators security group on each node. Additionally, for Windows 2003 the same cluster service account should be used for both the source and target.

5. Click **Next** to continue.
6. Choose the type of workload that you want to protect. Under **Server Workloads**, in the **Workload types** pane, select **Exchange Server**. In the **Workload items** pane, Double-Take will automatically select all of the Exchange databases and data files.

   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.

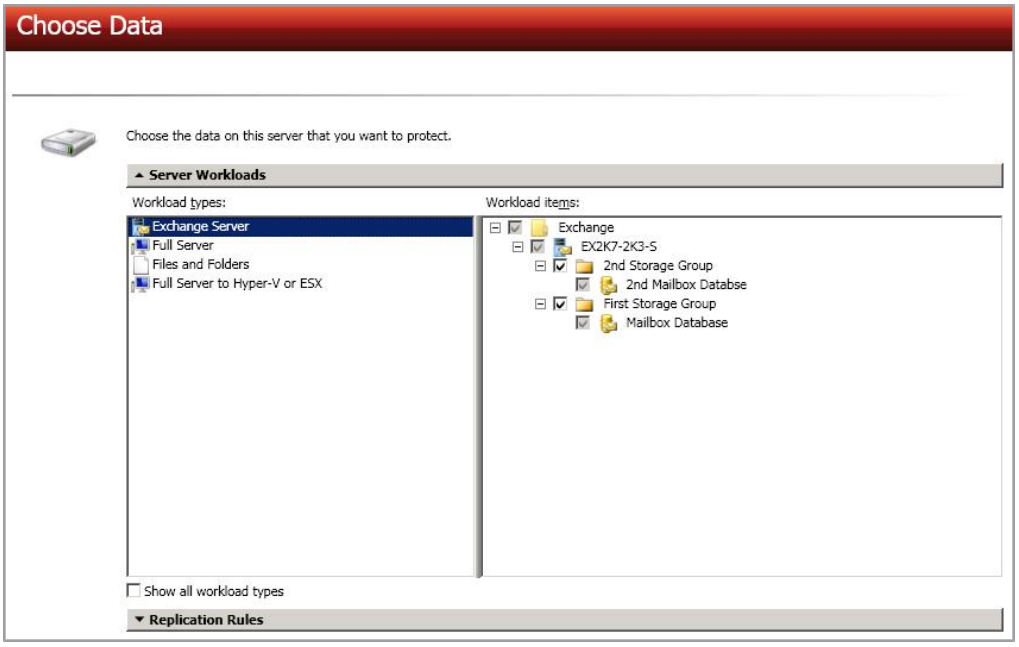

7. If you want to select other files and folders to include in your protection, click the **Replication Rules** heading and expand the volumes under **Folders**. For example, if you need to protect data that is stored on a non-mailbox server role, for example SMTP queue data, you will need to configure protection for that data separately.

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.

---

8. Click **Next** to continue.
9. Choose your target server. This is the server that will store the replica Exchange Server from the source.

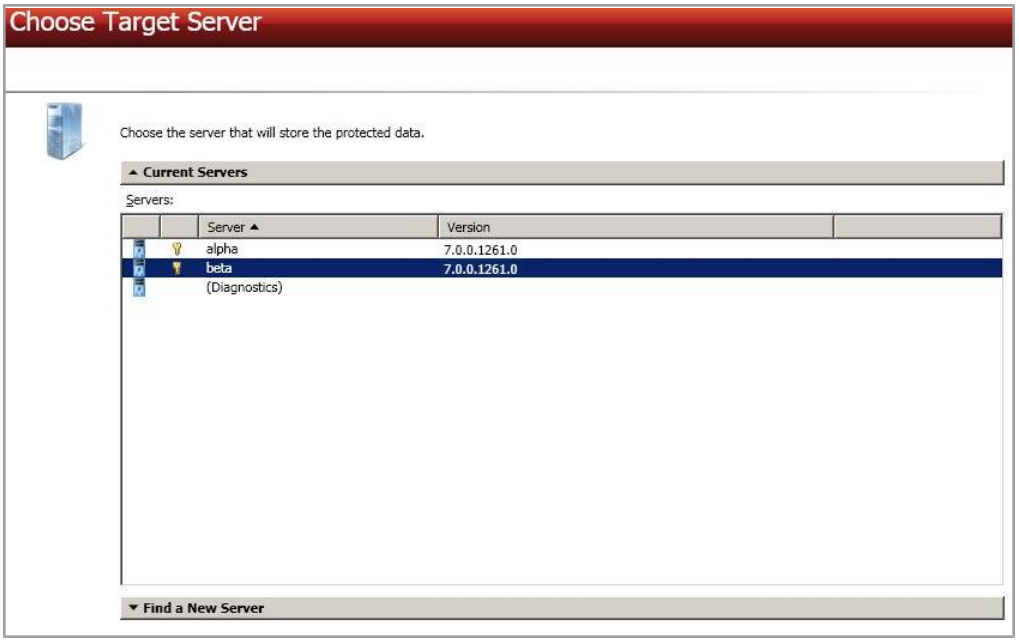

- **Current Servers**—This list contains the servers currently available in your console session. Servers that are not licensed for the workflow you have selected and those not applicable to the workload type you have selected will be filtered out of the list. Select your target server from the list.
- **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. If desired, the user must also meet the following credentials requirements, however, you can specify a separate user later in the job creation process to handle the Exchange specific tasks.

- The account must be an Exchange Full Administrator at the organizational level, as delegated via the Exchange System Manager at the user level.
- The account must have rights to manage Exchange in order to query and modify the Exchange Active Directory objects.
- The account must be a member of the Domain Admins group. If your security policies do not allow use of this group, see *DNS* on page 816 and use the instructions under the Double-Take DFO utility to use a non-Domain Admins account.
- If Exchange is on a cluster, the account must be a member of the Cluster Administrators security group on each node. Additionally, for Windows 2003 the same cluster service account should be used for both the source and target.

---

10. Click **Next** to continue.
11. You have many options available for your Exchange 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 320.

    - *General* on page 299
    - *Failover Monitor* on page 300
    - *Failover Options* on page 303
    - *Failover Identity* on page 305
    - *Mirror, Verify & Orphaned Files* on page 307
    - *Network Route* on page 311
    - *Failover Services* on page 312
    - *Snapshots* on page 313
    - *Compression* on page 314
    - *Bandwidth* on page 315

- *Scripts* on page 317
- *Test Failover Scripts* on page 319
- *Exchange Options* on page 320

General

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

Failover Monitor

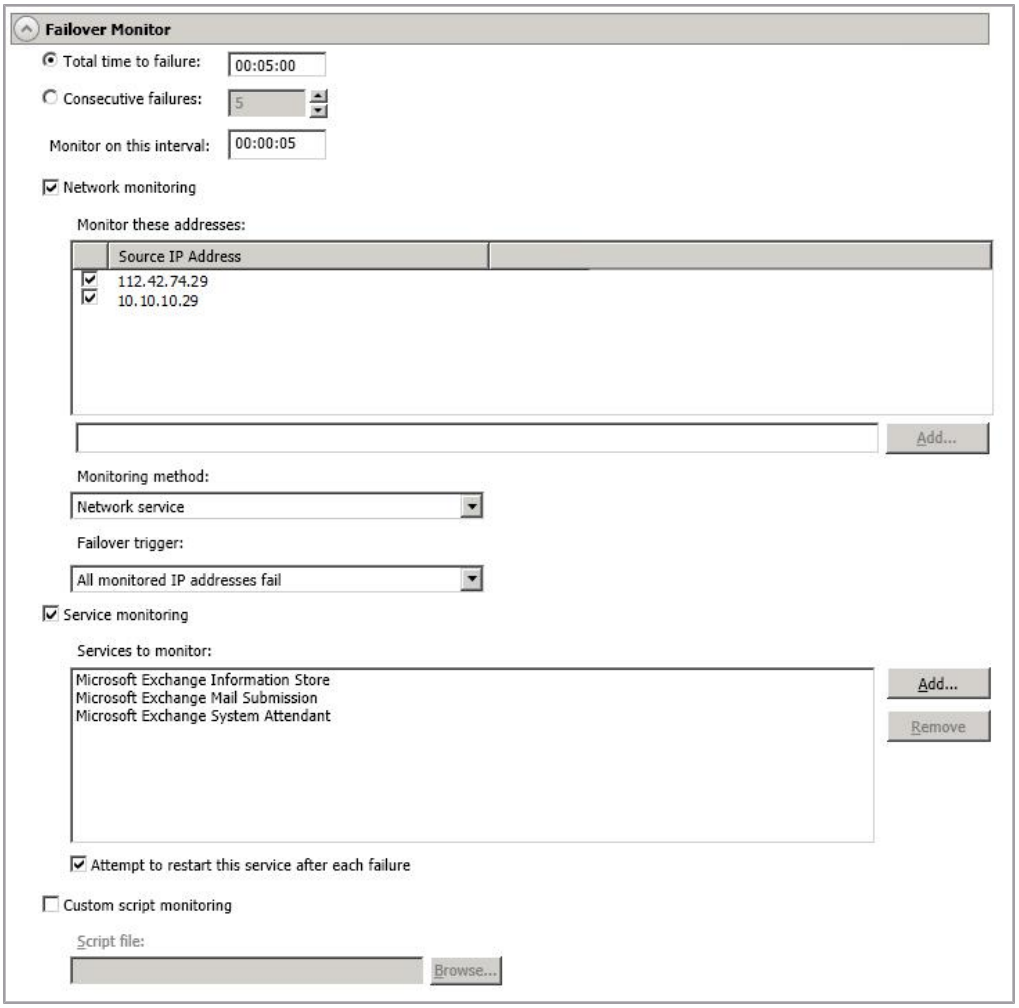

- **Total time to failure**—Specify, in hours:minutes:seconds, how long the target will keep trying to contact the source before the source is considered failed. This time is precise. If the total time has expired without a successful response from the source, this will be considered a failure.

  Consider a shorter amount of time for servers, such as a web server or order processing database, which must remain available and responsive at all times. Shorter times should be used where redundant interfaces and high-speed, reliable network links are available to prevent the false detection of failure. If the hardware does not support reliable communications, shorter times can lead to premature failover. Consider a longer amount of time for machines on slower networks or on a server that is not transaction critical. For example, failover would not be necessary in the case of a server restart.

- **Consecutive failures**—Specify how many attempts the target will make to contact the source before the source is considered failed. For example, if you have this option set to 20, and your source fails to respond to the target 20 times in a row , this will be considered a failure.
- **Monitor on this interval**—Specify, in hours:minutes:seconds, how long to wait between

attempts to contact the source to confirm it is online. This means that after a response (success or failure) is received from the source, Double-Take will wait the specified interval time before contacting the source again. If you set the interval to 00:00:00, then a new check will be initiated immediately after the response is received.

If you choose **Total time to failure**, do not specify a longer interval than failure time or your server will be considered failed during the interval period.

If you choose **Consecutive failures**, your failure time is calculated by the length of time it takes your source to respond plus the interval time between each response, times the number of consecutive failures that can be allowed. That would be (response time + interval) * failure number. Keep in mind that timeouts from a failed check are included in the response time, so your failure time will not be precise.

- **Network monitoring**—With this option, the target will monitor the source using a network ping.
  - **Monitor these addresses**—Select each **Source IP Address** that you want the target to monitor. If you want to monitor additional addresses, enter the address and click **Add**.

    

    If you are protecting a cluster, you are limited to the IP addresses in the cluster group that you are protecting.

  - **Monitoring method**—This option determines the type of network ping used for failover monitoring.
    - **Network service**—Source availability will be tested by an ICMP ping to confirm the route is active.
    - **Replication service**—Source availability will be tested by a UDP ping to confirm the Double-Take service is active.
    - **Network and replication services**—Source availability will be tested by both an ICMP ping to confirm the route is active and a UDP ping to confirm the Double-Take service is active. Both pings must fail in order to trigger a failover.
  - **Failover trigger**—If you are monitoring multiple IP addresses, specify when you want a failover condition to be triggered.
    - **One monitored IP address fails**—A failover condition will be triggered when any one of the monitored IP addresses fails. If each IP address is on a different subnet, you may want to trigger failover after one fails.
    - **All monitored IP addresses fail**—A failover condition will be triggered when all monitored IP addresses fail. If there are multiple, redundant paths to a server, losing one probably means an isolated network problem and you should wait for all IP addresses to fail.
- **Service monitoring**—This option is only available in standalone environments. It is not available in cluster environments.With this option, the target will monitor specific services on the source by confirming that they are running. Multiple services in the list will be checked in parallel. A failover condition is met when one of the monitored services fails the check. Click **Add** and select the service that you want to monitor. Repeat this step for additional services that you want to monitor. If you want to remove a service from the

Services to monitor list, highlight it and click **Remove**.

- **Attempt to restart this service after each failure**—When this option is enabled, if a service fails the monitor check, Double-Take will attempt to restart it. During this restart period, Double-Take will not check the other services, to avoid any false failures while the one service is attempting to be restarted. If the service cannot be restarted, Double-Take will consider this a failure.

- **Custom script monitoring**—This option is only available in standalone environments. It is not available in cluster environments. With this option, you can use your own custom script to monitor the source for a failure. You can use any standard script such as PowerShell, VB, batch files, and so on. Your script must return a code of 0 upon success, unless you are using a PowerShell script, in which case you should being using an exit value instead of a return value. Any other code will indicate a failure condition has been met. Your script will not be interactive, so ensure that it does not require any user interaction to execute successfully.

The network monitoring test is performed independently, while the service and custom monitoring tests are performed in parallel. Therefore, if you are using network monitoring with service and/or custom monitoring, and the network monitor fails, the service and custom monitoring will be skipped. If the network monitor is successful, service or custom monitoring will be performed. Any failure between these two tests will be considered a failure, therefore, both tests must complete and pass for the test to be considered successful.

Failover Options

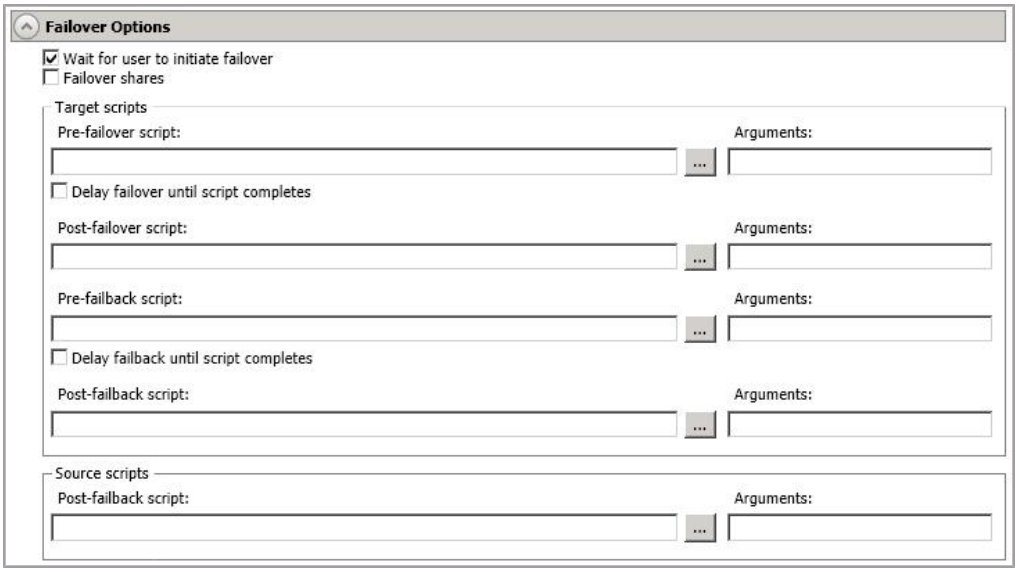

- **Wait for user to initiate failover**—By default, the failover process will wait for you to initiate it, allowing you to control when failover occurs. When a failure occurs, the job will wait in **Failover Condition Met** for you to manually initiate the failover process. Disable this option only if you want failover to occur immediately when a failure occurs.
- **Failover shares**—By default, shares will be failed over to the target. If you do not want to failover shares, disable this option.

  Automatic share failover only occurs for standard Windows file system shares. Other shares must be configured for failover through the failover scripts or created manually on the target. See *Macintosh shares* on page 826 or *NFS Shares* on page 827 for more information.

  If your target is a standalone server, Windows share information is automatically updated on the target once an hour. Since shares are not failed over, if your target is a cluster, you will need to manually update shares on a cluster target. For Windows 2003, you can create the file share resources while the source is running but keep the resource offline. For Windows 2008 and 2012, you will need to repeat the initial cluster configuration steps, although you can skip giving the target cluster account full control and creating the target Client Access Point, because those steps were already completed during the initial cluster configuration.

- **Scripts**—You can customize failover and failback by running scripts on the source and target. 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 99. Examples of functions specified in scripts include stopping services on the target before failover because they may not be necessary while the target is standing in for the source, stopping services on the target that need to be restarted with the source's machine

name and/or IP address, starting services or loading applications that are in an idle, standby mode waiting for failover to occur, notifying the administrator before and after failover or failback occurs, stopping services on the target after failback because they are no longer needed, stopping services on the target that need to be restarted with the target machine's original name and/or IP address, and so on. There are four types of failover and failback scripts.

- **Pre-failover script**—This script runs on the target at the beginning of the failover process. Specify the full path and name of the script file.
- **Post-failover script**—This script runs on the target at the end of the failover process. Specify the full path and name of the script file.
- **Pre-failback script** —This script runs on the target at the beginning of the failback process. Specify the full path and name of the script file.
- **Post-failback script**—This script runs on the target or source at the end of the failback process. Specify the full path and name of the script file.
- **Arguments**—Specify a comma-separated list of valid arguments required to execute the script.
- **Delay until script completes**—Enable this option if you want to delay the failover or failback process until the associated script has completed. If you select this option, make sure your script handles errors, otherwise the failover or failback process may never complete if the process is waiting on a script that cannot complete.

Scripts will run but will not be displayed on the screen if the Double-Take service is not set to interact with the desktop. Enable this option through the Windows Services applet.

Failover Identity

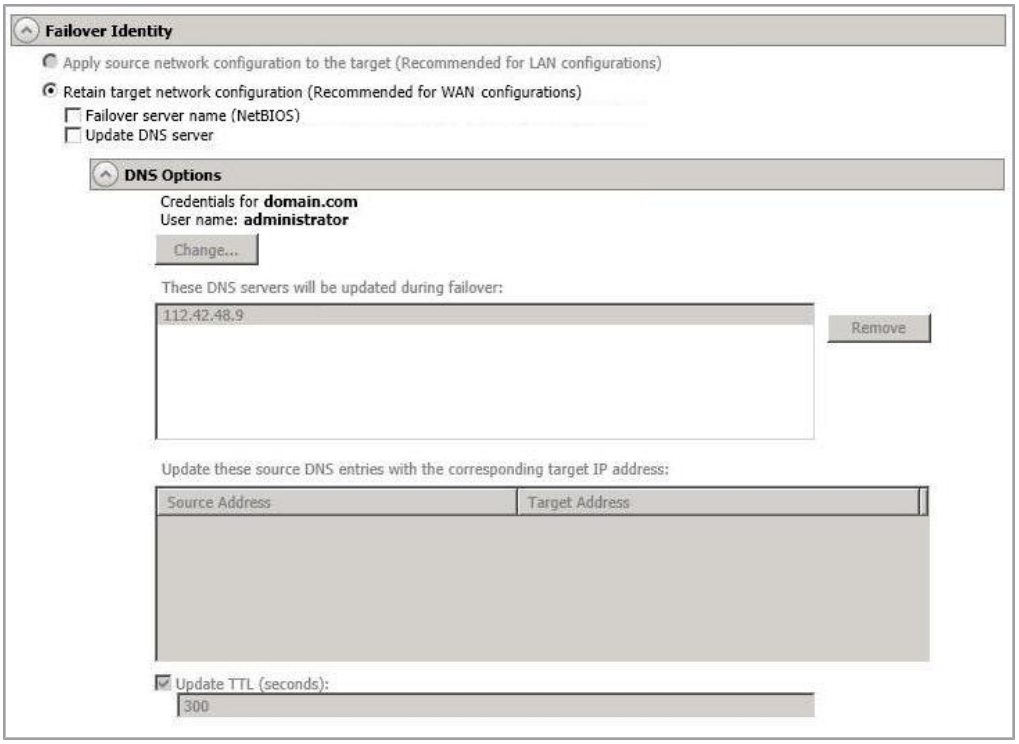

This section is not applicable to DAG to standalone configurations.

- **Retain target network configuration**—The target will retain all of its original IP addresses.
  - **Failover server name**—Select this option if you want to failover the NetBIOS name. This option will not be available for clusters.
  - **Update DNS server**—Specify if you want Double-Take to update your DNS server on failover. If DNS updates are made, the DNS records will be locked during failover. Be sure and review the *Core Double-Take requirements* on page 23 for the requirements for updating DNS.

    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. For clusters, be sure and select the clustered IP address.
- **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.

If you select **Retain your target network configuration** but do not enable **Update DNS server**, you will need to specify failover scripts that update your DNS server during failover, or you can update the DNS server manually after failover. This would also apply to non--Microsoft Active Directory Integrated DNS servers. You will want to keep your target network configuration but do not update DNS. In this case, you will need to specify failover scripts that update your DNS server during failover, or you can update the DNS server manually after failover.

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

Mirror, Verify & Orphaned Files

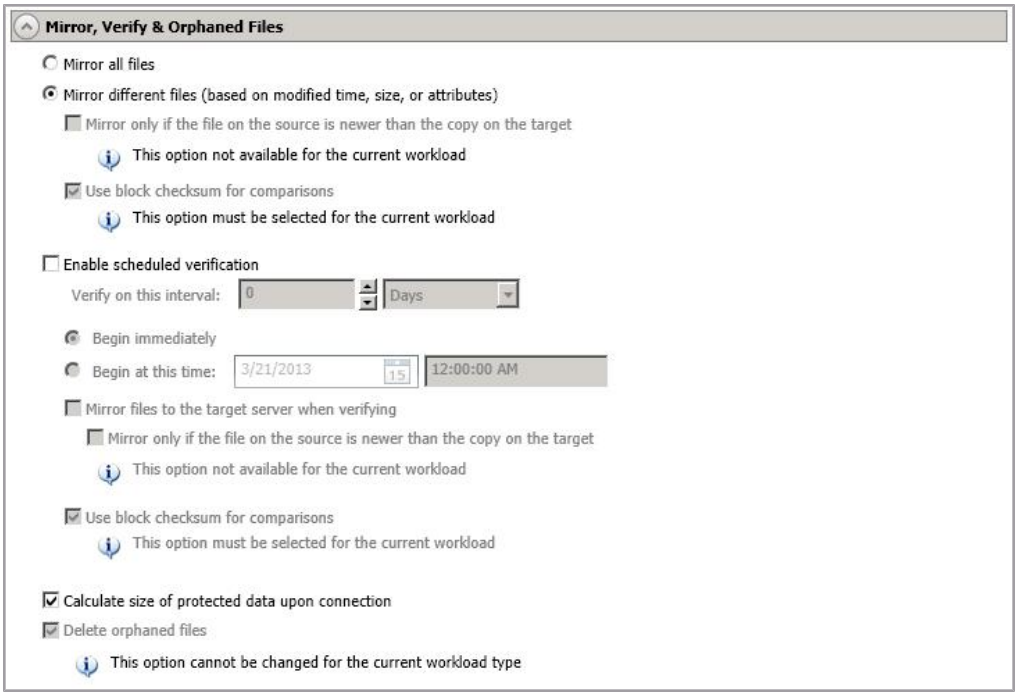

- **Mirror all files**—All protected files will be mirrored from the source to the target.
- **Mirror different files**—Only those protected files that are different based on date and time, size, or attributes will be mirrored from the source to the target.
  - **Mirror only if the file on the source is newer than the copy on the target**—This option is not available for Exchange jobs.
  - **Use block checksum for comparisons**—For those files flagged as different, the mirroring process can perform a block checksum comparison and send only those blocks that are different. This option cannot be disabled for Exchange jobs.

File Differences Mirror Options Compared

The following table will help you understand how the various difference mirror options work together, including when you are using the block checksum option configured through the *Source server properties* on page 92.

An X in the table indicates that option is enabled. An X enclosed in parentheses (X) indicates that the option can be on or off without impacting the action performed during the mirror.

Not all job types have the source newer option available.

Source Server Properties	Job Properties			Action Performed
Block Checksum Option	File Differences Option	Source Newer Option	Block Checksum Option	
(X)	X			Any file that is different on the source and target based on the date, time, size, and/or attribute is transmitted to the target. The mirror sends the entire file.
(X)	X	X		Any file that is newer on the source than on the target based on date and/or time is transmitted to the target. The mirror sends the entire file.
	X		X	Any file that is different on the source and target based on date, time, size, and/or attributed is flagged as different. The mirror then performs a checksum comparison on the flagged files and only sends those blocks that are different.
X	X		X	The mirror performs a checksum comparison on all files and only sends those blocks that are different.
(X)	X	X	X	Any file that is newer on the source than on the target based on date and/or time is flagged as different. The mirror then performs a checksum comparison on the flagged files and only sends those blocks that are different.

- **Enable scheduled verification**—Verification is the process of confirming that the source replica data on the target is identical to the original 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, can automatically initiate a remirror, if configured. The remirror ensures data integrity between the source and target. When this option is enabled, Double-Take will verify the source replica data on the target and generate a verification log.

  

  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.

  - **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 at the specified date and time.
  - **Mirror files to the target server when verifying**—When this option is enabled, in addition to verifying the data and generating a log, Double-Take will also mirror to the target any protected files that are different on the source.
    - **Mirror only if the file on the source is newer than the copy on the target**—This option is not available for Exchange jobs.
    - **Use block checksum for comparisons**—For those files flagged as different, the mirroring process can perform a block checksum comparison and send only those blocks that are different.
- **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 the calculation is enabled, it is completed before the job starts mirroring, which can take a significant amount of time depending on the number of files and system performance. If your job contains a large number of files, for example, 250,000 or more, you may want to disable the calculation so that data will start being mirrored sooner. 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.

- **Delete orphaned files**—An orphaned file is a file that exists in the replica data on the target, but does not exist in the protected data on the source. This option specifies if orphaned files should be deleted on the target during a mirror, verification, or restoration. This option cannot be disabled for Exchange jobs because it is important that you delete orphaned files because if you have orphaned Exchange log files on the target, it is possible

that one or more Exchange database will not mount after failover.

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

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 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.

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 95 for more information.

During a mirror, orphaned file processing success messages will be logged to a separate orphaned file log. 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 100 for details on the location of that folder. The orphaned log file is overwritten during each orphaned file processing during a mirror, and the log file will be a maximum of 50 MB.

Network Route

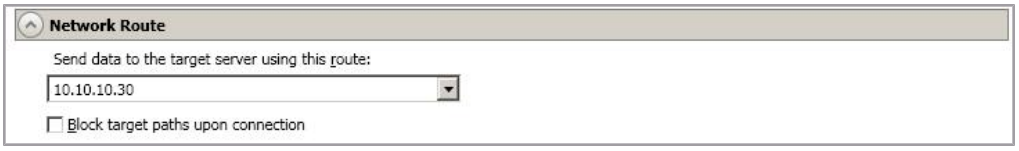

- **Send data to the target server using this route**—By default, Double-Take will select a target route for transmissions. If desired, specify an alternate route on the target that the data will be transmitted through. This allows you to select a different route for Double-Take traffic. For example, you can separate regular network traffic and Double-Take traffic on a machine with multiple IP addresses.

  The IP address used on the source will be determined through the Windows route table.

  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.

- **Block target paths upon 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. Any target paths that are blocked will be unblocked automatically during the failover process so that users can modify data after failover. During restoration, the paths are automatically blocked again. If you failover and failback without performing a restoration, the target paths will remain unblocked.

Failover Services

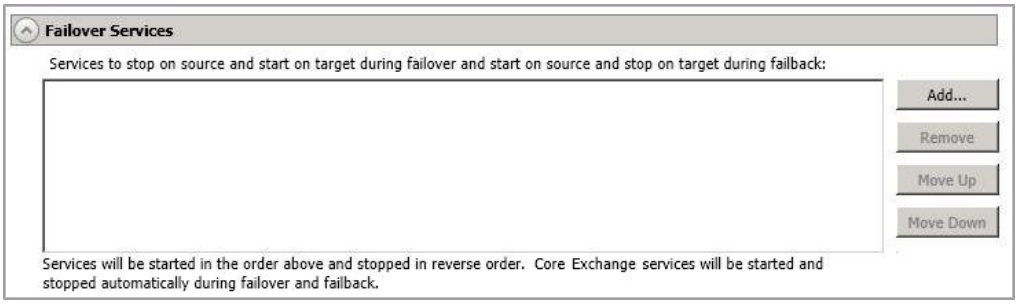

This section is not applicable to clustered environments.

Services to stop on source and start on target during failover and start on source and stop on target during failback—Double-Take automatically determines the appropriate Exchange services to start and stop based on your Exchange version, Exchange configuration, and your operating system. If necessary, you can start and stop other services during failover and failback. Click **Add** to insert a service into the list or **Remove** to remove a service from the list. The services will be started in the order they appear and stopped in the reverse order. Highlight a service and click **Move Up** or **Move Down** to arrange the services in the desired order.

Snapshots

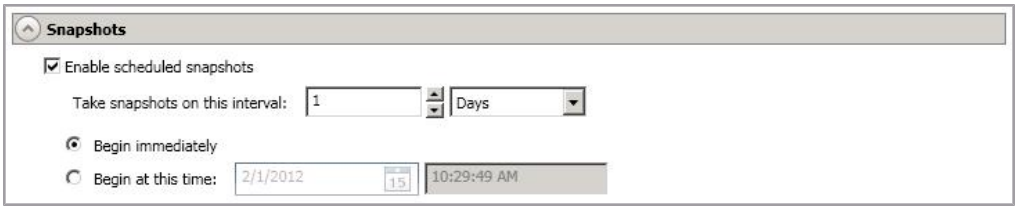

This section is not applicable to clustered environments.

A snapshot is an image of the source replica data on the target taken at a single point in time. You can view the snapshots in VSS and recover any files or folders desired. You can also failover to a snapshot.

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 starting at a later date and time. Specify the date and time parameters to indicate when you want to start.

See *Managing snapshots* on page 161 for details on taking manual snapshots and deleting snapshots.

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

Compression

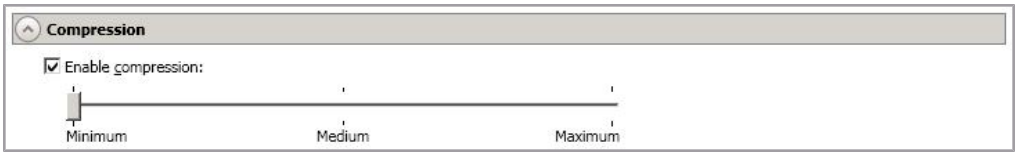

To help reduce the amount of bandwidth needed to transmit Double-Take data, compression allows you to compress data prior to transmitting it across the network. In a WAN environment this provides optimal use of your network resources. If compression is enabled, the data is compressed before it is transmitted 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

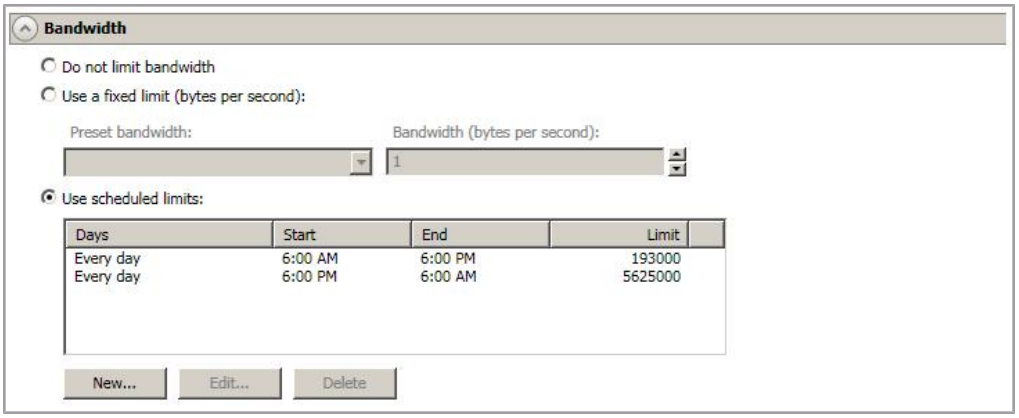

Bandwidth limitations are available to restrict the amount of network bandwidth used for Double-Take data transmissions. When a bandwidth limit is specified, Double-Take never exceeds that allotted amount. The bandwidth not in use by Double-Take is available for all other network traffic.

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

The scheduled option is not available if your source is a cluster.

- **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

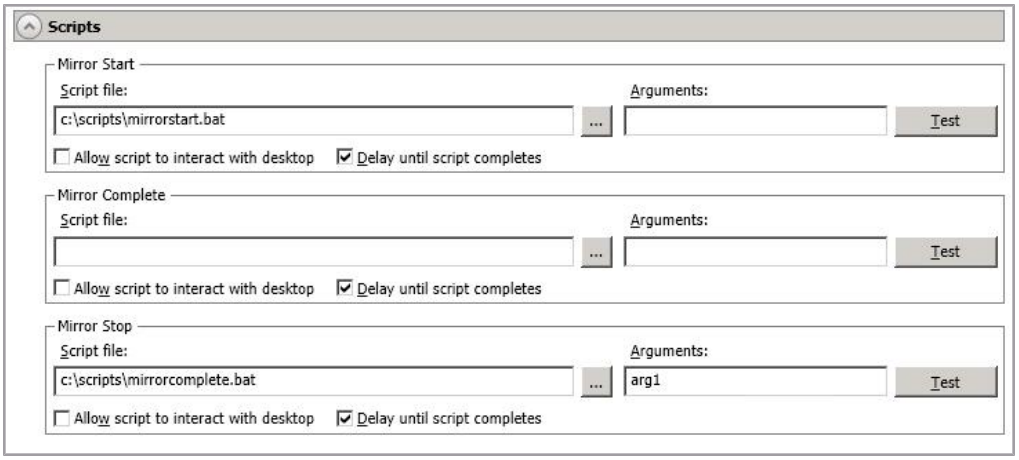

You can customize mirroring by running scripts on the target at pre-defined points in the mirroring process. Scripts may contain any valid Windows command, executable, or batch file. The scripts are processed using the same account running the Double-Take service, unless you have identified a specific account through the server's properties. See *Script credentials* on page 99. 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.
- **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.

Mirror scripts are dependent on the target and the **Target Path Mappings** specified under the **Network Route & Folder Selection** section. 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).

Test Failover Scripts

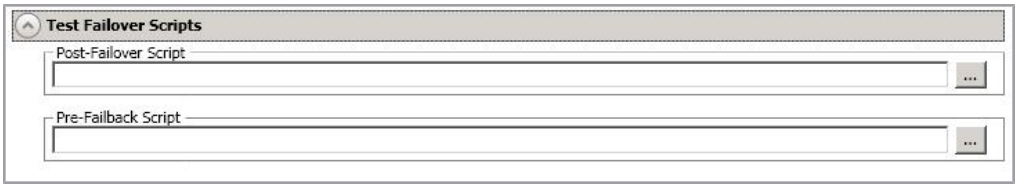

This section is not applicable to clustered environments.

When you failover, you will have three choices. You can failover to live data, failover to data from a snapshot, or perform a test failover. Any scripts you specify in this section are only used during a test failover.

- **Post-Failover Script**—This script runs on the target at the end of the failover process. Specify the full path and name of the script file.
- **Pre-Failback Script** —This script runs on the target at the beginning of the undo failover process. Specify the full path and name of the script file.

Scripts will run but will not be displayed on the screen if the Double-Take service is not set to interact with the desktop. Enable this option through the Windows Services applet.

Exchange Options

Click **Change** and specify a user that meets the following requirements.

- The login account must be an Exchange Full Administrator at the organizational level, as delegated via the Exchange System Manager at the user level.
- The login account must have rights to manage Exchange in order to query and modify the Exchange Active Directory objects.
- If Exchange is on a cluster, the login account must be a member of the Cluster Administrators security group on each node. Additionally, for Windows 2003 the same cluster service account should be used for both the source and target.

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

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

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

Managing and controlling Exchange jobs

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

The jobs displayed in the right pane depend on the server group folder selected in the left pane. Every job for each server in your console session is displayed when the **Jobs on All Servers** group is selected. If you have created and populated server groups (see *Managing servers* on page 65), then only the jobs associated with the server or target servers in that server group will be displayed in the right pane.

- See *Overview job information displayed in the top pane* on page 321
- See *Detailed job information displayed in the bottom pane* on page 323
- See *Job controls* on page 325

Overview job information displayed in the top pane

The top pane displays high-level overview information about your jobs.

Column 1 (Blank)

The first blank column indicates the state of the job.

 The job is in a healthy state.

 The job is in a warning state. This icon is also displayed on any server groups that you have created that contain a job in a warning state.

 The job is in an error state. This icon is also displayed on any server groups that you have created that contain a job in an error state.

 The job is in an unknown state.

Job

The name of the job

Source Server

The name of the source. This could be a name or IP address of a standalone server, a cluster, or a node. Cluster jobs will be associated with the cluster name and standalone jobs will be associated with a standalone server or a cluster node.

Target Server

The name of the target. This could be a name or IP address of a standalone server, a cluster, or a node. Cluster jobs will be associated with the cluster name and standalone jobs will be associated with a standalone server or a cluster node.

Job Type

Each job type has a unique job type name. This job is an Exchange Server job. For a complete list of all job type names, press F1 to view the Double-Take Console online help.

Activity

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

Mirror Status

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

Replication Status

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

Transmit Mode

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

Detailed job information displayed in the bottom pane

The details displayed in the bottom pane of the **Manage Jobs** page provide additional information for the job highlighted in the top pane. If you select multiple jobs, the details for the first selected job will be displayed.

Name

The name of the job

Target data state

- **OK**—The data on the target is in a good state.
- **Mirroring**—The target is in the middle of a mirror process. The data will not be in a good state until the mirror is complete.
- **Mirror Required**—The data on the target is not in a good state because a remirror is required. This may be caused by an incomplete or stopped mirror or an operation may have been dropped on the target.
- **Restore Required**—The data on the source and target do not match because of a failover condition. Restore the data from the target back to the source. If you want to discard the changes on the target, you can remirror to resynchronize the source and target.
- **Snapshot Reverted**—The data on the source and target do not match because a snapshot has been applied on the target. Restore the data from the target back to the source. If you want to discard the changes on the target, you can remirror to resynchronize the source and target.
- **Busy**—The source is low on memory causing a delay in getting the state of the data on the target.
- **Not Loaded**—Double-Take target functionality is not loaded on the target server. This may be caused by an activation code error.
- **Unknown**—The console cannot determine the status.

Mirror remaining

The total number of mirror bytes that are remaining to be sent from the source to the target

Mirror skipped

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

Replication queue

The total number of replication bytes in the source queue

Disk queue

The amount of disk space being used to 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**.

Connected since

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

Recent activity

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

Additional information

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

Job controls

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

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

Create a New Job 

This button leaves the **Manage Jobs** page and opens the **Get Started** page.

View Job Details 

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

Delete 

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 77. You will remain on the **Manage Jobs** page after updating the server credentials. If your servers use the same credentials, make sure you also update the credentials on the **Manage Servers** page so that the Double-Take Console can authenticate to the servers in the console session. See *Managing servers* on page 65.

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. Failover monitoring will continue 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. Failover monitoring will continue while the job is stopped. 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.

Snapshots are not applicable to clustered environments.

Manage Snapshots 

Allows you to manage your snapshots by taking and deleting snapshots for the selected job. See *Managing snapshots* on page 161 for more information.

Snapshots are not applicable to clustered environments.

Failover or Cutover 

Starts the failover process. See *Failing over Exchange jobs* on page 339 for the process and details of failing over an Exchange job.

Failback 

Starts the failback process. See *Restoring then failing back Exchange jobs* on page 341 for the process and details of failing back an Exchange job.

Restore 

Starts the restoration process. See *Restoring then failing back Exchange jobs* on page 341 for the process and details of restoring an Exchange job.

Reverse

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

Undo Failover

Cancels a test failover by undoing it. Undo failover does not apply to clustered jobs. See *Failing over Exchange jobs* on page 339 for details on undoing a test failover.

View Job Log

Opens the job log. On the right-click menu, this option is called **View Logs**, and you have the option of opening the job log, source server log, or target server log. See *Viewing the log files through the Double-Take Console* on page 678 for details on all three of these logs.

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 all files**—All protected files will be mirrored from the source to the target.
  - **Mirror different files**—Only those protected files that are different based on date and time, size, or attributes will be mirrored from the source to the target.
    - **Mirror only if the file on the source is newer than the copy on the target**—Only those protected files that are newer on the source are mirrored to the target.

    

    If you are using a database application or are protecting a domain controller, do not use this option 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 that might be newer, get mirrored.

- **Use block checksum for comparisons**—For those files flagged as different, the mirroring process can perform a block checksum comparison and send only those blocks that are different.

- **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 the calculation is enabled, it is completed before the job starts mirroring, which can take a significant amount of time depending on the number of files and system performance. If your job contains a large number of files, for example, 250,000 or more, you may want to disable the calculation so that data will start being mirrored sooner. 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.
  - **Create verification report only**—This option verifies the data and generates a verification log, but it does not remirror any files that are different on the source and target. See *Verification log* on page 102 for details on the log file.
  - **Mirror files to the target server automatically**—When this option is enabled, in addition to verifying the data and generating a log, Double-Take will also mirror to the target any protected files that are different on the source.
    - **Mirror only if the file on the source is newer than the copy on the target**—Only those protected files that are newer on the source are mirrored to the target.

      

      If you are using a database application or are protecting a domain controller, do not use this option 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 that might be newer, get mirrored.

    - **Use block checksum for comparisons**—For those files flagged as different, the mirroring process can perform a block checksum comparison and send only those blocks that are different.
- **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.
- **Update Shares**—Windows share information is automatically updated on the target once an hour. This option allows you to manually update share information immediately when the option is selected. Shares are not applicable to environments where the target is a cluster.

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 65.

Type a server name

Displays only jobs that contain the text you entered. If you have created and populated server groups, then only jobs that contain the text you entered associated with the server or target servers in that server group will be displayed. See *Managing servers* on page 65.

Overflow Chevron 

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

Viewing Exchange job details

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

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

Job name

The name of the job

Job type

Each job type has a unique job type name. This job is an Exchange Server job. For a complete list of all job type names, press F1 to view the Double-Take Console online help.

Health

 The job is in a healthy state.

 The job is in a warning state.

 The job is in an error state.

 The job is in an unknown state.

Activity

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

Connection ID

The incremental counter used to number connections. The number is incremented when a connection is created. It is also incremented by internal actions, such as an auto-disconnect and auto-reconnect. The lowest available number (as connections are created, stopped, deleted, and so on) will always be used. The counter is reset to one each time the Double-Take service is restarted.

Transmit mode

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

Target data state

- **OK**—The data on the target is in a good state.
- **Mirroring**—The target is in the middle of a mirror process. The data will not be in a good state until the mirror is complete.
- **Mirror Required**—The data on the target is not in a good state because a remirror is required. This may be caused by an incomplete or stopped mirror or an operation may have been dropped on the target.
- **Restore Required**—The data on the source and target do not match because of a failover condition. Restore the data from the target back to the source. If you want to discard the changes on the target, you can remirror to resynchronize the source and target.
- **Snapshot Reverted**—The data on the source and target do not match because a snapshot has been applied on the target. Restore the data from the target back to the source. If you want to discard the changes on the target, you can remirror to resynchronize the source and target.
- **Busy**—The source is low on memory causing a delay in getting the state of the data on the target.
- **Not Loaded**—Double-Take target functionality is not loaded on the target server. This may be caused by an activation code error.
- **Unknown**—The console cannot determine the status.

Target route

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

Compression

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

Encryption

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

Bandwidth limit

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

Connected since

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

Additional information

Depending on the current state of your job, you may see additional information displayed to keep you informed about the progress and status of your job. If there is no

additional information, you will see (None) displayed.

Mirror status

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

Mirror percent complete

The percentage of the mirror that has been completed

Mirror remaining

The total number of mirror bytes that are remaining to be sent from the source to the target

Mirror skipped

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

Replication status

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

Replication queue

The total number of replication bytes in the source queue

Disk queue

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

Bytes sent

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

Bytes sent compressed

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

Validating an Exchange job

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

1. From the **Manage Jobs** page, highlight the job and click **View Job Details** in the toolbar.
2. In the **Tasks** area on the right on the **View Job Details** page, click **Validate job properties**.
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 enable protection. Depending on the error, you may be able to click **Fix** or **Fix All** and let Double-Take correct the problem for you. For those errors that Double-Take cannot correct automatically, you will need to modify the source or target to correct the error, or you can select a different target. You must revalidate the selected servers, by clicking **Recheck**, until the validation check passes without errors.

   Validation checks for an existing job are logged to the job log on the target server. See *Log files* on page 677 for details on the various log files.
4. Once your servers have passed validation, click **Close**.

Editing an Exchange job

Use these instructions to edit an Exchange job.

1. From the **Manage Jobs** page, highlight the job and click **View Job Details** in the toolbar.
2. In the **Tasks** area on the right on the **View Job Details** page, click **Edit job properties**. (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 have the same options available for your Exchange job as when you created the job. See *Creating an Exchange job* on page 295 for details on each job option.

   

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

4. If you want to modify the workload items or replication rules for the job, click **Edit workload or replication rules**. Modify the **Workload item** you are protecting, if desired. Additionally, you can modify the specific **Replication Rules** for your job.

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

   You can also enter wildcard rules, however you should do so carefully. Rules are applied to files that are closest in the directory tree to them. If you have rules that include multiple folders, an exclusion rule with a wild card will need to be added for each folder that it needs applied to. For example, if you want to exclude all .log files from D:\ and your rules include D:\, D:\Dir1 , and D:\Dir2, you would need to add the exclusion rule for the root and each subfolder rule. So you will need to add exclude rules for D:*.log , D:\Dir1*.log, and D:\Dir2*.log.

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

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

   Errors are designated by a white X inside a red circle. Warnings are designated by a black exclamation point (!) inside a yellow triangle. A successful validation is designated by a white checkmark inside a green circle. You can sort the list by the icon to see errors, warnings, or successful validations together. Click on any of the validation items to see details. You must correct any errors before you can enable protection. Depending on the error, you may be able to

click **Fix** or **Fix All** and let Double-Take correct the problem for you. For those errors that Double-Take cannot correct automatically, you will need to modify the source or target to correct the error, or you can select a different target. You must revalidate the selected servers, by clicking **Recheck**, until the validation check passes without errors.

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.

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

Viewing an Exchange job log

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

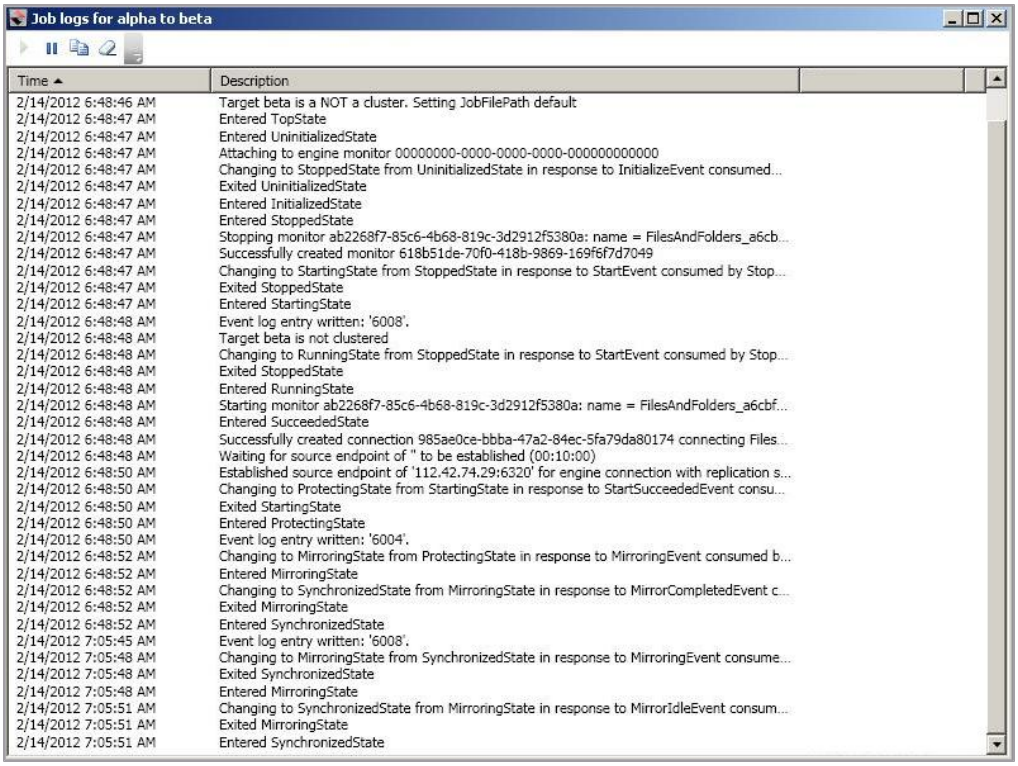

The following table identifies the controls and the table columns in the **Job logs** window.

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.

Failing over Exchange jobs

When a failover condition has been met, failover will be triggered automatically if you disabled the wait for user option during your failover configuration. If the wait for user before failover option is enabled, you will be notified in the console when a failover condition has been met. At that time, you will need to trigger it manually from the console when you are ready.

1. On the **Manage Jobs** page, highlight the job that you want to failover and click **Failover or Cutover** in the toolbar.
2. Select the type of failover to perform.
   - **Failover to live data**—Select this option to initiate a full, live failover using the current data on the target. The application services will be stopped on the source (if it is online), and they will be started on the target. DNS records will be updated, if applicable. Application requests destined for the source server or its IP addresses are routed to the target.
   - **Perform test failover**—Select this option to perform a test failover using the current data on the target for standalone environments. This option is not applicable to clustered environments.This option will allow you to confirm Exchange on the target is viable for failover. This process will pause the target (the entire target is not paused, just Double-Take processing), take a snapshot of the target (so that you can revert back to the pre-test state after the test is complete), and then start the application services on the target. Success and failure messages will be available in the job log. (Note the application services are still running on the source during the test, and not users are moved during the test.) Once the application services are running on the target, you can perform any testing desired on the target server.

     

     The following caveats apply to performing a test failover for Exchange.
     - Make sure you have enough space on your target to contain the snapshot.
     - If you have scheduled Double-Take snapshots, they will continue to be created during the test, however, you should not use any of these snapshots from this time period because the target data may not be in a good state.
     - If you are using Windows 2008 R2 and Double-Take snapshots, any snapshots that you take before testing failover will be unusable for an actual failover. To work around this issue, take another snapshot immediately after the test. Once the additional, post-testing snapshot is taken, all of the previous snapshots will be usable for future failovers. This issue is fixed in the Windows 2008 R2 Service Pack 1 release. If you install that release or a later release, you will not have to worry about the extra snapshot after the test.
     - If you have taken multiple Double-Take snapshots on the target and then tested failover, in some instances the SMTPSVC service may not start after failing over to a snapshot. If possible, failover to a snapshot taken after the target data was tested. If that is not possible, you will have to start the SMTPSVC service manually on the target after failover.

- If you are protecting Exchange 2007 or Exchange 2010 and need to use the Exchange Management Console while the Double-Take Availability test failover process is running, you will need to start the IIS service. Stop the IIS service after the test failover is complete.

- **Failover to a snapshot**—Select this option to initiate a full, live failover without using the current data on the target. Instead, select a snapshot and the data on the target will be reverted to that snapshot. This option will not be available if there are no snapshots on the target or if the target does not support snapshots. This option is also not applicable to clustered environments. To help you understand what snapshots are available, the **Type** indicates the kind of snapshot.
  - **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. Select how you want to handle the data in the target queue. You may want to check the amount of data in queue on the target by reviewing the *Statistics* on page 688 or *Performance Monitor* on page 790.
   - **Apply data in target queues before failover or cutover**—All of the data in the target queue will be applied before failover begins. The advantage to this option is that all of the data that the target has received will be applied before failover begins. 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 failover or cutover immediately**—All of the data in the target queue will be discarded and failover will begin immediately. The advantage to this option is that failover will occur immediately. The disadvantage is that any data in the target queue will be lost.
   - **Revert to last good snapshot if target data state is bad**—If the target data is in a bad state, Double-Take will automatically revert to the last good Double-Take snapshot before failover begins. If the target data is in a good state, Double-Take will not revert the target data. Instead, Double-Take will apply the data in the target queue and then failover. The advantage to this option is that good data on the target is guaranteed to be used. The disadvantage is that if the target data state is bad, you will lose any data between the last good snapshot and the failure.
4. When you are ready to begin failover, click **Failover**.
5. If you performed a test failover, you can undo it by selecting **Undo Failover or Cutover** in the toolbar. The application services will be stopped on the target, the pre-test snapshot that was taken will be reverted, and the target will be resumed.

Restoring then failing back Exchange jobs

Restoring before failing back allows your users to continue accessing their data on the failed over target, which is standing in for the source, while you perform the restoration process.

1. Resolve the problem(s) on the source that caused it to fail.
2. Bring the source onto the network.

In a consolidated Exchange environment, you may not be able to login to the source using domain credentials, while the source is failed over.

---

3. On the **Manage Jobs** page, highlight the job and click **Restore**.

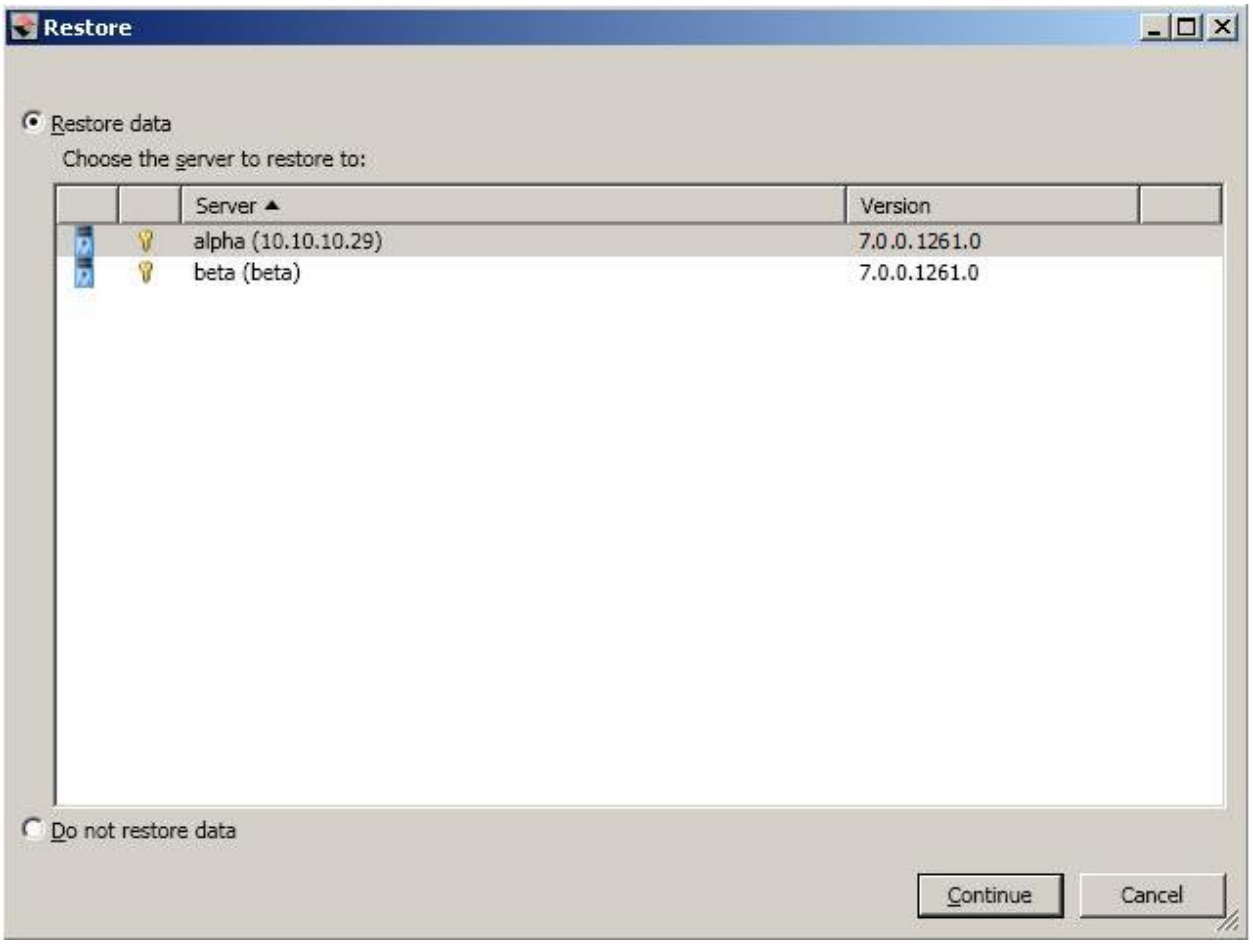

4. Confirm **Restore data** is selected, then highlight your source server in the server list.

If you do not want to restore your data, you can select **Do not restore data**. Keep in mind, that any data changes that took place on the target after failover will be lost.

---

5. Click **Continue** to start the restoration.
6. During the restoration process, the **Activity** will indicate **Restoring** and **Mirror Status** will indicate **In Progress**. When the restoration is complete, the **Mirror Status** will change to **Idle**, and the **Activity** will be **Restored**. At this time, schedule a time for failback. User downtime will begin once failback is started, so select a time that will have minimal disruption on your users.
7. On the **Manage Jobs** page, highlight the job and click **Failback**.
8. In the dialog box, highlight the job that you want to failback and click **Failback**.
9. Check that your source is fully functional and that the source data is in a good state, and then, if desired, you can enable protection again by clicking **Start**.

Chapter 10 SQL protection

Create a SQL job when you have Microsoft SQL Server and want application-level protection.

- See *SQL requirements* on page 344—SQL protection includes specific requirements for this type of protection.
- See *Creating a SQL job* on page 347—This section includes step-by-step instructions for creating a SQL job.
- See *Managing and controlling SQL jobs* on page 372—You can view status information about your SQL jobs and learn how to control these jobs.
- See *Failing over SQL jobs* on page 390—Use this section when a failover condition has been met or if you want to failover manually.
- *Restoring then failing back SQL jobs* on page 392—Use this section when you are ready to restore and failback.

If your source is a domain controller, you should use one of the full server protection methods to protect the entire server as a best practice. See *Selecting a protection type* on page 165.

SQL requirements

In addition to the *Core Double-Take requirements* on page 23, you must also meet the following requirements to protect SQL.

- **SQL versions**—Double-Take can protect Microsoft SQL Server or Express 2005, 2008, 2008 R2, or 2012.
- **SQL and network configuration**—The following requirements and limitations apply to your SQL server and network configuration.
  - All Microsoft best practices should be used for all versions of SQL.
  - You should use the same version, service pack, and architecture (32-bit or 64-bit) of SQL Server on both the source and target servers.
  - SQL 2012 is the only supported SQL version if you are using Windows 2012 in a cluster configuration.
  - If you are using SQL 2008 R2, cluster to cluster configurations are supported, however, the AlwaysOn cluster feature is not supported.
  - If you have multiple SQL instances on the same cluster, each instance must be on a separate node before you can create a job. If the instances are ever on the same node, you will need to delete the job, move the instances back to separate nodes, and then re-create the job. This is the only method that will guarantee data integrity.
  - If you are using SQL Express 2008, you will need to enable and start the SQL Browser Service. By default, this service is set to disabled. Additionally, you will need to enable TCP/IP to accept remote connections. To do this, launch the SQL Server Configuration Manager. Expand **SQL Server Network Configuration**, and under **Protocols** enable **TCP/IP**.
  - If you are using SQL Express 2005, you will need to enable named pipes and TCP/IP to accept remote connections. To do this, launch the SQL Server Configuration Manager. Expand **SQL Server 2005 Network Configuration**, and under **Protocols**enable **Named Pipes** and **TCP/IP**.
  - The SQL program files must be installed in the same location on the source and target.
  - The drive letter(s) where SQL stores its data on the source must be the same on the target.
  - The source and target servers must have named instances with the same name installed prior to configuring protection.
  - Single-label DNS domain names (those without a suffix such as .com, .corp, .net) are not supported.
  - In environments where the FIPS security policy is enabled, you must use impersonation, which requires the following.
    - The user running the Double-Take Console must have all appropriate rights to update the domain (that is, only impersonation is supported).
    - You must manually verify DNS rights by running the DFO utility with the /test parameter.
  - Microsoft Server Core is not supported.
  - If your source and target are in a domain, they should be in the same domain. If they are not, the SQL Server service on both the source and target servers must be configured to start with the same domain user account.

- If your source and target are in a workgroup, make sure the source server's NIC does not register the IP addresses with DNS.
- If you are using a domain service account that is not in the domain or local Administrators security group, the replicated databases will not mount on the target after failover or on the source after restore because of security restrictions at the file system level. You need to place the SQL 2005 service account in the local Administrators group on the source and target because of the way these versions use local groups for NTFS permissions.
- You may want to exclude the tempdb database to reduce mirroring and replication traffic. See *Application optimizations* on page 836.
- Transparent Data Encryption (TDE) is supported for SQL 2008, however the SQL service must be running with the same service account on the source and target.
- You should use a domain user account as the Windows Service Account for SQL. See Setting Up Windows Service Accounts on the Microsoft web site for more information. You should include this account in the local Administrators group on your source and target to ensure that the appropriate permissions for the replicated databases and log files will be available after failover and failback.
- Double-Take does not support a SQL default instance that is using non-default ports.

- **Snapshots**—You can take and failover to Double-Take snapshots using a SQL job, however snapshots are not supported if your source and/or target is a cluster. See *Core Double-Take requirements* on page 23 for the specific snapshot requirements.
- **Supported configurations**—The following table identifies the supported configurations for a SQL job.

Configuration		Supported	Not Supported
Source to target configuration	One to one, active/standby	X	
	One to one, active/active		X
	Many to one		X
	One to many		X
	Chained		X
	Single server		X
Server configuration	Standalone to standalone	X	
	Standalone to cluster		X
	Cluster to standalone	X	
	Cluster to cluster	X	
	Cluster Shared Volumes (CSV) guest level	X	
	Cluster Shared Volumes (CSV) host level		X

Configuration		Supported	Not Supported
Upgrade configuration[1]	Upgrade 5.3 Double-Take Application Manager job to 7.0 Double-Take Console SQL job		X
	Upgrade 6.0 SQL job to 7.0 SQL job	X	
Version 7.0 console[2]	Version 7.0 console can create job for 5.3 source and 5.3 target		X
	Version 7.0 console can create job for 6.0 source and 6.0 target	X	
	Version 7.0 console can create job for 7.0 source and 7.0 target	X	

1. For an upgrade configuration that is not supported, you will have to delete the existing job before the upgrade and create a new job after the upgrade.
2. Newer job options available in the version 7.0 console will not be functional when creating jobs for servers running version 6.0.

Creating a SQL job

Use these instructions to create a SQL job.

If you have a multiple SQL instances on the same cluster, each instance must be on a separate node before you can create a job. If the instances are ever on the same node, you will need to delete the job, move the instances back to separate nodes, and then re-create the job. This is the only method that will guarantee data integrity.

---

1. Make sure you are logged in to the Double-Take Console as a user with the SQL sysadmin role before you begin the SQL job creation process.
2. Click **Get Started** from the toolbar.
3. Select **Double-Take Availability** and click **Next**.
4. Select **Protect files and folders, an application, or an entire Windows server** and click **Next**.
5. Choose your source server. This is the physical or virtual server running SQL Server. If your source is a cluster, select the cluster name, not the SQL virtual server name or virtual IP address.

- **Current Servers**—This list contains the servers currently available in your console session. Servers that are not licensed for the workflow you have selected will be filtered out of the list. Select your source server from the list.
- **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 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.

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. Additionally, the user must meet the following credentials requirements.

- The account must be assigned the sysadmin role on the SQL server in order to query and administer SQL, and the SQL service startup account should be a domain account.
- The account must be a member of the Domain Admins group. If your security policies do not allow use of this group, see *DNS* on page 816 and use the instructions under the Double-Take DFO utility to use a non-Domain Admins account.

6. Click **Next** to continue.
7. Choose the type of workload that you want to protect. Under **Server Workloads**, in the **Workload types** pane, select **SQL Server**. In the **Workload items** pane, Double-Take will automatically select the entire SQL program and data files. If desired, you can exclude user databases from protection, but the system databases (except for tempdb) are required and cannot be excluded.

   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.

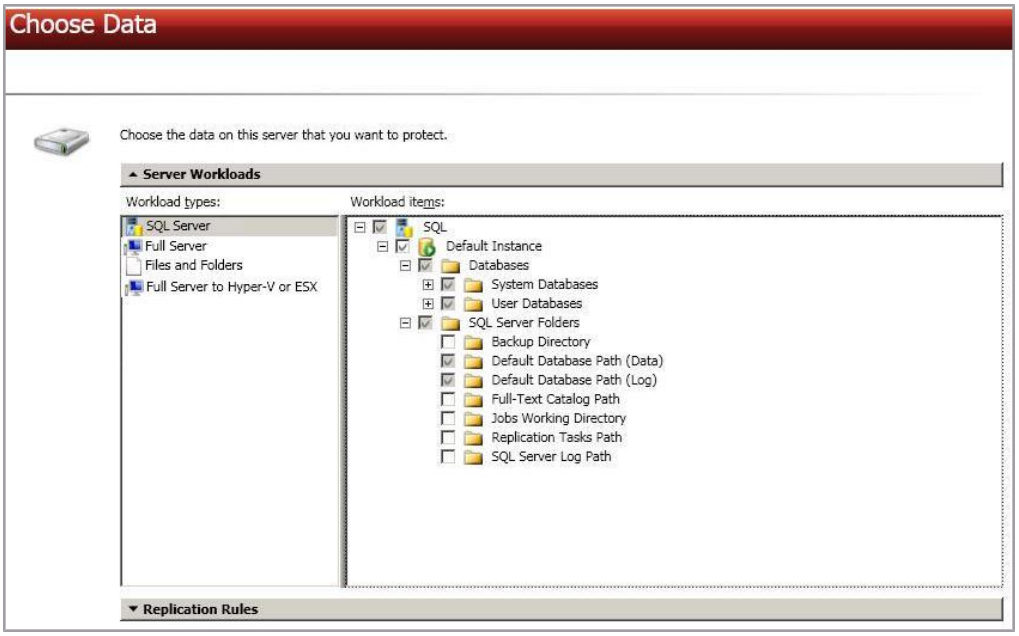

8. If you want to select other files and folders to include in your protection, click the **Replication Rules** heading and expand the volumes under **Folders**.

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.

---

9. Click **Next** to continue.
10. Choose your target server. This is the server that will store the replica SQL Server from the source.

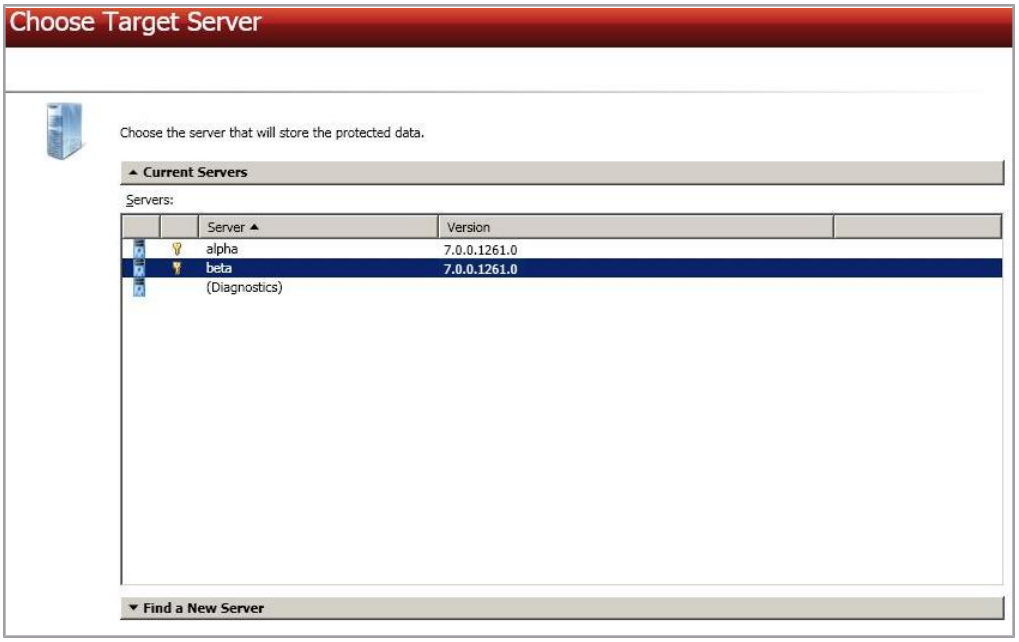

- **Current Servers**—This list contains the servers currently available in your console session. Servers that are not licensed for the workflow you have selected and those not applicable to the workload type you have selected will be filtered out of the list. Select your target server from the list.
- **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. Additionally, the user must meet the following credentials requirements.

- The account must be assigned the sysadmin role on the SQL server in order to query and administer SQL, and the SQL service startup account should be a domain account.
- The account must be a member of the Domain Admins group. If your security policies do not allow use of this group, see *DNS* on page 816 and use the instructions under the Double-Take DFO utility to use a non-Domain Admins account.

---

11. Click **Next** to continue.
12. You have many options available for your SQL 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 371.

    - *General* on page 351
    - *Failover Monitor* on page 352
    - *Failover Options* on page 355
    - *Failover Identity* on page 357
    - *Mirror, Verify & Orphaned Files* on page 359
    - *Network Route* on page 363
    - *Failover Services* on page 364
    - *Snapshots* on page 365
    - *Compression* on page 366
    - *Bandwidth* on page 367
    - *Scripts* on page 369
    - *Test Failover Scripts* on page 371

General

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

Failover Monitor

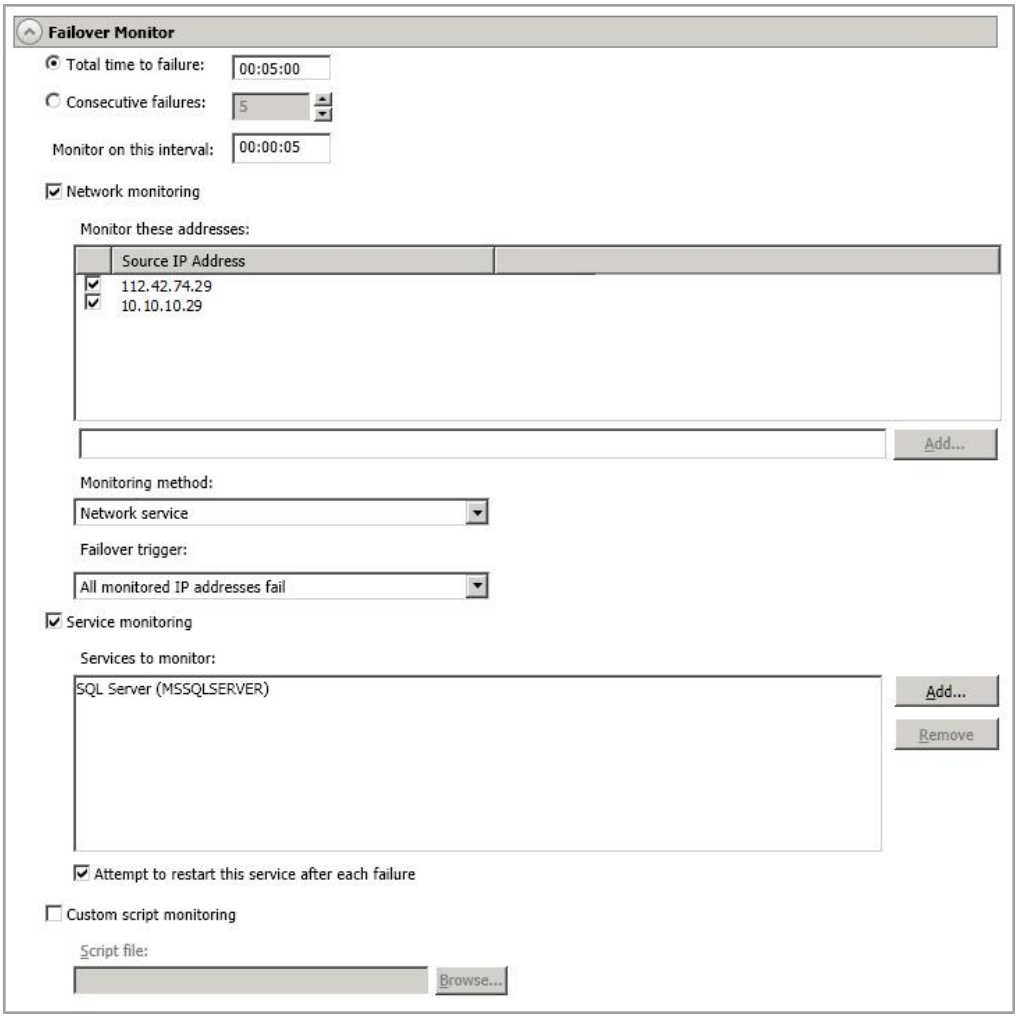

- **Total time to failure**—Specify, in hours:minutes:seconds, how long the target will keep trying to contact the source before the source is considered failed. This time is precise. If the total time has expired without a successful response from the source, this will be considered a failure.

  Consider a shorter amount of time for servers, such as a web server or order processing database, which must remain available and responsive at all times. Shorter times should be used where redundant interfaces and high-speed, reliable network links are available to prevent the false detection of failure. If the hardware does not support reliable communications, shorter times can lead to premature failover. Consider a longer amount of time for machines on slower networks or on a server that is not transaction critical. For example, failover would not be necessary in the case of a server restart.

- **Consecutive failures**—Specify how many attempts the target will make to contact the source before the source is considered failed. For example, if you have this option set to 20, and your source fails to respond to the target 20 times in a row , this will be considered a failure.
- **Monitor on this interval**—Specify, in hours:minutes:seconds, how long to wait between

attempts to contact the source to confirm it is online. This means that after a response (success or failure) is received from the source, Double-Take will wait the specified interval time before contacting the source again. If you set the interval to 00:00:00, then a new check will be initiated immediately after the response is received.

If you choose **Total time to failure**, do not specify a longer interval than failure time or your server will be considered failed during the interval period.

If you choose **Consecutive failures**, your failure time is calculated by the length of time it takes your source to respond plus the interval time between each response, times the number of consecutive failures that can be allowed. That would be (response time + interval) * failure number. Keep in mind that timeouts from a failed check are included in the response time, so your failure time will not be precise.

- **Network monitoring**—With this option, the target will monitor the source using a network ping.
  - **Monitor these addresses**—Select each **Source IP Address** that you want the target to monitor. If you want to monitor additional addresses, enter the address and click **Add**.

    

    If you are protecting a cluster, you are limited to the IP addresses in the cluster group that you are protecting.

  - **Monitoring method**—This option determines the type of network ping used for failover monitoring.
    - **Network service**—Source availability will be tested by an ICMP ping to confirm the route is active.
    - **Replication service**—Source availability will be tested by a UDP ping to confirm the Double-Take service is active.
    - **Network and replication services**—Source availability will be tested by both an ICMP ping to confirm the route is active and a UDP ping to confirm the Double-Take service is active. Both pings must fail in order to trigger a failover.
  - **Failover trigger**—If you are monitoring multiple IP addresses, specify when you want a failover condition to be triggered.
    - **One monitored IP address fails**—A failover condition will be triggered when any one of the monitored IP addresses fails. If each IP address is on a different subnet, you may want to trigger failover after one fails.
    - **All monitored IP addresses fail**—A failover condition will be triggered when all monitored IP addresses fail. If there are multiple, redundant paths to a server, losing one probably means an isolated network problem and you should wait for all IP addresses to fail.
- **Service monitoring**—This option is only available in standalone environments. It is not available for clustered source servers. With this option, the target will monitor specific services on the source by confirming that they are running. Multiple services in the list will be checked in parallel. A failover condition is met when one of the monitored services fails the check. Click **Add** and select the service that you want to monitor. Repeat this step for additional services that you want to monitor. If you want to remove a service from the

Services to monitor list, highlight it and click **Remove**.

- **Attempt to restart this service after each failure**—When this option is enabled, if a service fails the monitor check, Double-Take will attempt to restart it. During this restart period, Double-Take will not check the other services, to avoid any false failures while the one service is attempting to be restarted. If the service cannot be restarted, Double-Take will consider this a failure.

- **Custom script monitoring**—This option is only available in standalone environments. It is not available for clustered source servers. With this option, you can use your own custom script to monitor the source for a failure. You can use any standard script such as PowerShell, VB, batch files, and so on. Your script must return a code of 0 upon success, unless you are using a PowerShell script, in which case you should being using an exit value instead of a return value. Any other code will indicate a failure condition has been met. Your script will not be interactive, so ensure that it does not require any user interaction to execute successfully.

The network monitoring test is performed independently, while the service and custom monitoring tests are performed in parallel. Therefore, if you are using network monitoring with service and/or custom monitoring, and the network monitor fails, the service and custom monitoring will be skipped. If the network monitor is successful, service or custom monitoring will be performed. Any failure between these two tests will be considered a failure, therefore, both tests must complete and pass for the test to be considered successful.

Failover Options

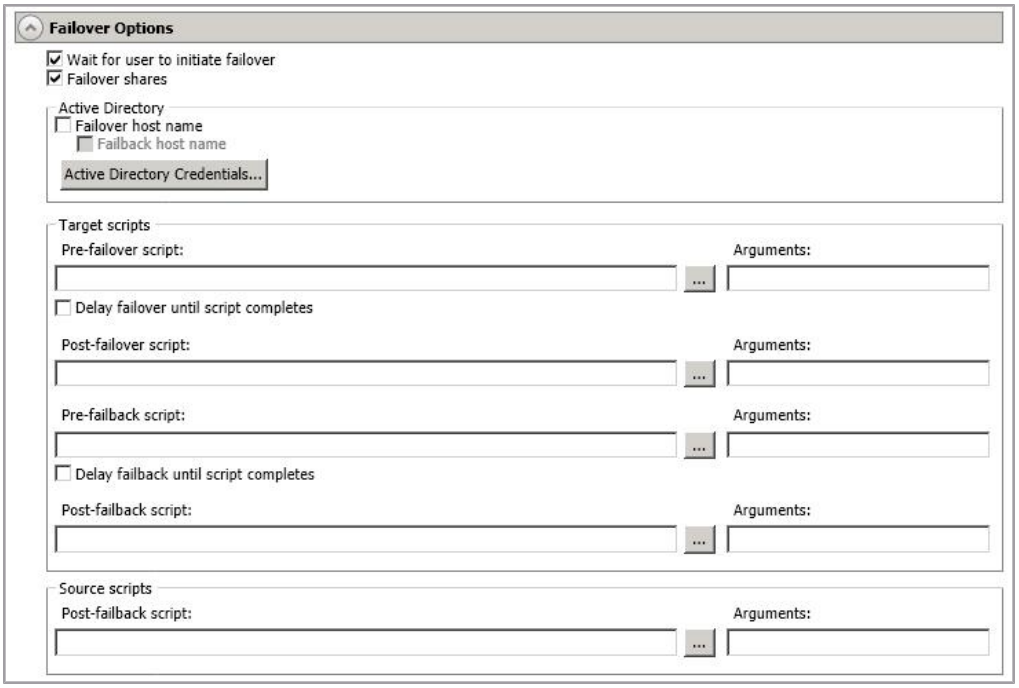

- **Wait for user to initiate failover**—By default, the failover process will wait for you to initiate it, allowing you to control when failover occurs. When a failure occurs, the job will wait in **Failover Condition Met** for you to manually initiate the failover process. Disable this option only if you want failover to occur immediately when a failure occurs.
- **Failover shares**—By default, shares will be failed over to the target. If you do not want to failover shares, disable this option.

  Automatic share failover only occurs for standard Windows file system shares. Other shares must be configured for failover through the failover scripts or created manually on the target. See *Macintosh shares* on page 826 or *NFS Shares* on page 827 for more information.

  If your target is a standalone server, Windows share information is automatically updated on the target once an hour. Since shares are not failed over, if your target is a cluster, you will need to manually update shares on a cluster target. For Windows 2003, you can create the file share resources while the source is running but keep the resource offline. For Windows 2008 and 2012, you will need to repeat the initial cluster configuration steps, although you can skip giving the target cluster account full control and creating the target Client Access Point, because those steps were already completed during the initial cluster configuration.

- **Failover host name**—If desired, you can failover the source server's host name. This will automatically remove the host SPN (Service Principle Name) from Active Directory on the source and add it to Active Directory on the target. If you are using two different accounts for the SQL service login on the source and target, you should failover the SPNs. If you are

using the same domain account for the SQL service login on the source and target, you do not need to failover the SPNs.

- **Failback host name**—This option returns the host SPN on the source and target back to their original settings on failback. If you are using Active Directory, enable this option or you may experience problems with failback.
- **Active Directory Credentials**—If you are failing over and/or failing back the host name, you need to specify a user that has update privileges within Active Directory. Click **Active Directory Credentials** and identify a user and the associated password that has privileges to create and delete SPNs. The username must be in the format fully_qualified_ domain\user,, and the account password cannot be blank.
- **Scripts**—You can customize failover and failback by running scripts on the source and target. 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 99. Examples of functions specified in scripts include stopping services on the target before failover because they may not be necessary while the target is standing in for the source, stopping services on the target that need to be restarted with the source's machine name and/or IP address, starting services or loading applications that are in an idle, standby mode waiting for failover to occur, notifying the administrator before and after failover or failback occurs, stopping services on the target after failback because they are no longer needed, stopping services on the target that need to be restarted with the target machine's original name and/or IP address, and so on. There are four types of failover and failback scripts.
  - **Pre-failover script**—This script runs on the target at the beginning of the failover process. Specify the full path and name of the script file.
  - **Post-failover script**—This script runs on the target at the end of the failover process. Specify the full path and name of the script file.
  - **Pre-failback script** —This script runs on the target at the beginning of the failback process. Specify the full path and name of the script file.
  - **Post-failback script**—This script runs on the target or source at the end of the failback process. Specify the full path and name of the script file.
  - **Arguments**—Specify a comma-separated list of valid arguments required to execute the script.
  - **Delay until script completes**—Enable this option if you want to delay the failover or failback process until the associated script has completed. If you select this option, make sure your script handles errors, otherwise the failover or failback process may never complete if the process is waiting on a script that cannot complete.

  Scripts will run but will not be displayed on the screen if the Double-Take service is not set to interact with the desktop. Enable this option through the Windows Services applet.

Failover Identity

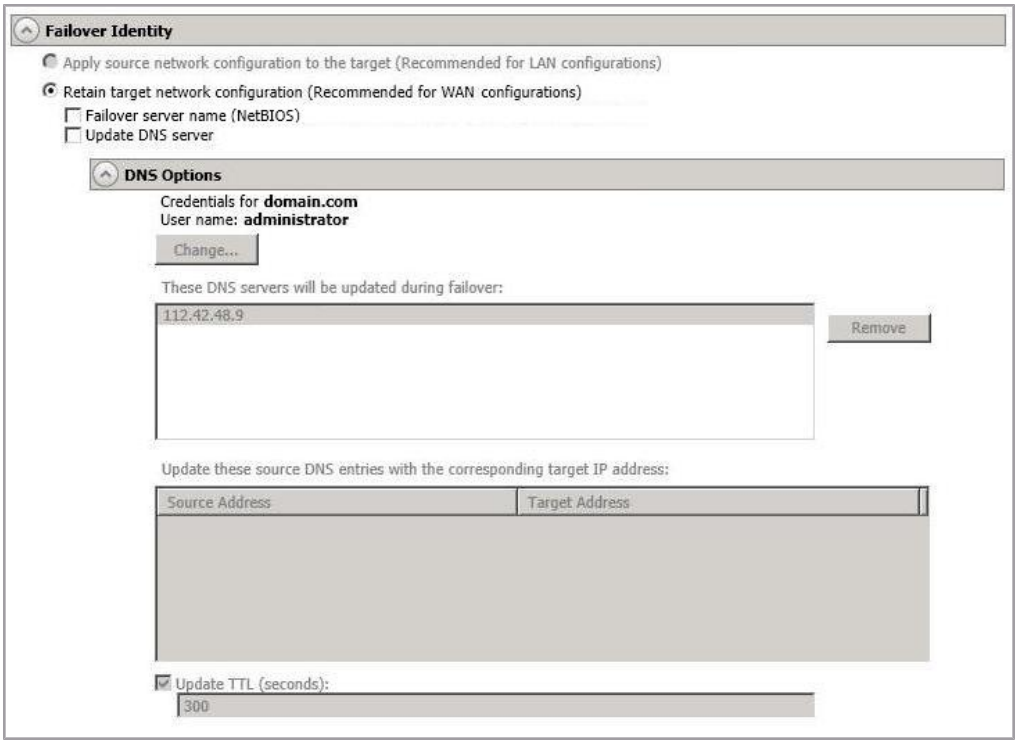

- **Retain target network configuration**—The target will retain all of its original IP addresses.
  - **Failover server name**—Select this option if you want to failover the NetBIOS name. This option will not be available for clusters.
  - **Update DNS server**—Specify if you want Double-Take to update your DNS server on failover. If DNS updates are made, the DNS records will be locked during failover. Be sure and review the *Core Double-Take requirements* on page 23 for the requirements for updating DNS.

    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. For clusters, be sure and select the clustered IP address.

- **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.

If you select **Retain your target network configuration** but do not enable **Update DNS server**, you will need to specify failover scripts that update your DNS server during failover, or you can update the DNS server manually after failover. This would also apply to non--Microsoft Active Directory Integrated DNS servers. You will want to keep your target network configuration but do not update DNS. In this case, you will need to specify failover scripts that update your DNS server during failover, or you can update the DNS server manually after failover.

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

Mirror, Verify & Orphaned Files

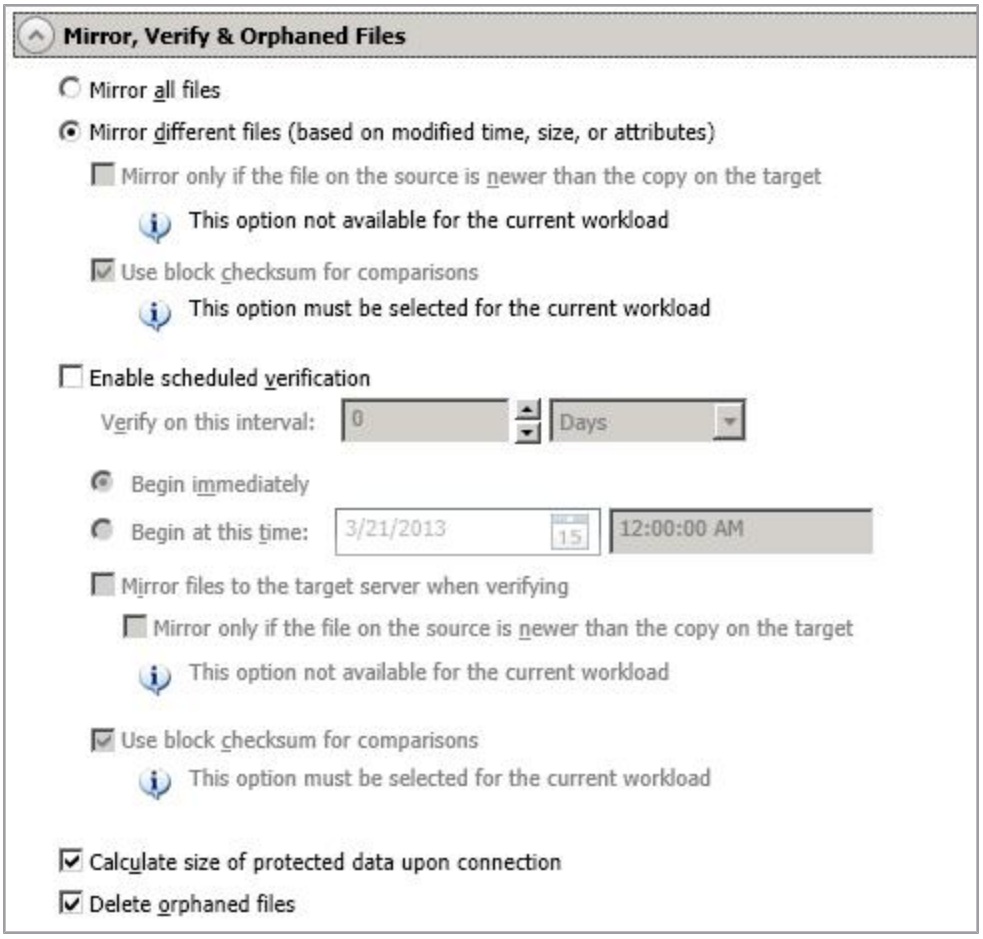

- **Mirror all files**—All protected files will be mirrored from the source to the target.
- **Mirror different files**—Only those protected files that are different based on date and time, size, or attributes will be mirrored from the source to the target.
  - **Mirror only if the file on the source is newer than the copy on the target**—This option is not available for SQL jobs.
  - **Use block checksum for comparisons**—For those files flagged as different, the mirroring process can perform a block checksum comparison and send only those blocks that are different. This option cannot be disabled for SQL jobs.

File Differences Mirror Options Compared

The following table will help you understand how the various difference mirror options work together, including when you are using the block checksum option configured through the *Source server properties* on page 92.

An X in the table indicates that option is enabled. An X enclosed in parentheses (X) indicates that the option can be on or off without impacting the action performed during the mirror.

Not all job types have the source newer option available.

Source Server Properties	Job Properties			Action Performed
Block Checksum Option	File Differences Option	Source Newer Option	Block Checksum Option	
(X)	X			Any file that is different on the source and target based on the date, time, size, and/or attribute is transmitted to the target. The mirror sends the entire file.
(X)	X	X		Any file that is newer on the source than on the target based on date and/or time is transmitted to the target. The mirror sends the entire file.
	X		X	Any file that is different on the source and target based on date, time, size, and/or attributed is flagged as different. The mirror then performs a checksum comparison on the flagged files and only sends those blocks that are different.
X	X		X	The mirror performs a checksum comparison on all files and only sends those blocks that are different.
(X)	X	X	X	Any file that is newer on the source than on the target based on date and/or time is flagged as different. The mirror then performs a checksum comparison on the flagged files and only sends those blocks that are different.

- **Enable scheduled verification**—Verification is the process of confirming that the source replica data on the target is identical to the original 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, can automatically initiate a remirror, if configured. The remirror ensures data integrity between the source and target. When this option is enabled, Double-Take will verify the source replica data on the target and generate a verification log.

  

  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.

  - **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 at the specified date and time.
  - **Mirror files to the target server when verifying**—When this option is enabled, in addition to verifying the data and generating a log, Double-Take will also mirror to the target any protected files that are different on the source.
    - **Mirror only if the file on the source is newer than the copy on the target**—This option is not available for SQL jobs.
    - **Use block checksum for comparisons**—For those files flagged as different, the mirroring process can perform a block checksum comparison and send only those blocks that are different.
- **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 the calculation is enabled, it is completed before the job starts mirroring, which can take a significant amount of time depending on the number of files and system performance. If your job contains a large number of files, for example, 250,000 or more, you may want to disable the calculation so that data will start being mirrored sooner. 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.

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

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

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 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.

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 95 for more information.

During a mirror, orphaned file processing success messages will be logged to a separate orphaned file log. 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 100 for details on the location of that folder. The orphaned log file is overwritten during each orphaned file processing during a mirror, and the log file will be a maximum of 50 MB.

Network Route

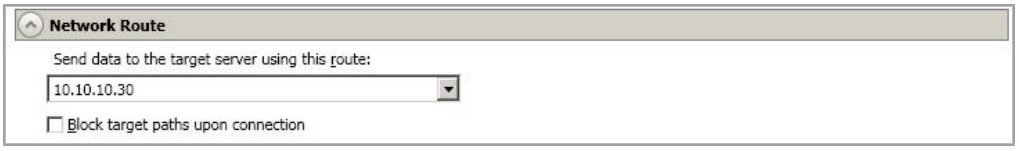

- **Send data to the target server using this route**—By default, Double-Take will select a target route for transmissions. If desired, specify an alternate route on the target that the data will be transmitted through. This allows you to select a different route for Double-Take traffic. For example, you can separate regular network traffic and Double-Take traffic on a machine with multiple IP addresses.

  The IP address used on the source will be determined through the Windows route table.

  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.

- **Block target paths upon 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. Any target paths that are blocked will be unblocked automatically during the failover process so that users can modify data after failover. During restoration, the paths are automatically blocked again. If you failover and failback without performing a restoration, the target paths will remain unblocked.

Failover Services

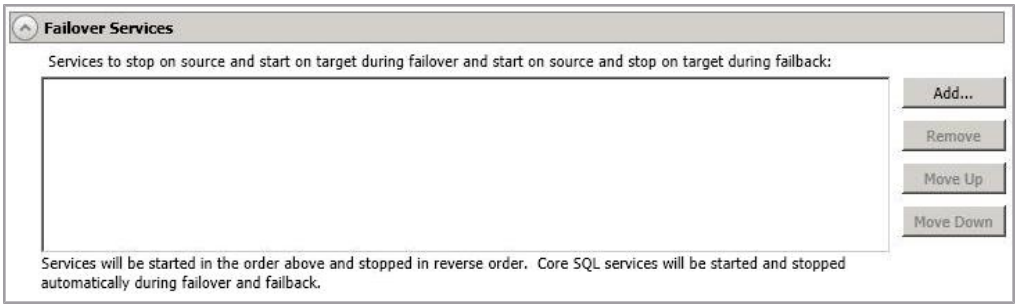

This section is not applicable to clustered environments.

Services to stop on source and start on target during failover and start on source and stop on target during failback—Double-Take automatically determines the appropriate SQL services to start and stop based on your SQL version, SQL configuration, and your operating system. If necessary, you can start and stop other services during failover and failback. Click **Add** to insert a service into the list or **Remove** to remove a service from the list. The services will be started in the order they appear and stopped in the reverse order. Highlight a service and click **Move Up** or **Move Down** to arrange the services in the desired order.

Snapshots

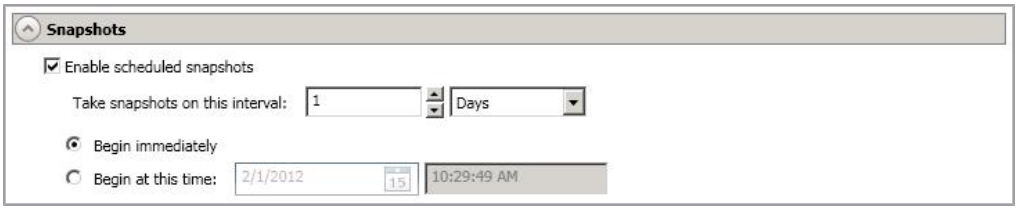

This section is not applicable to clustered environments.

A snapshot is an image of the source replica data on the target taken at a single point in time. You can view the snapshots in VSS and recover any files or folders desired. You can also failover to a snapshot.

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 starting at a later date and time. Specify the date and time parameters to indicate when you want to start.

See *Managing snapshots* on page 161 for details on taking manual snapshots and deleting snapshots.

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

Compression

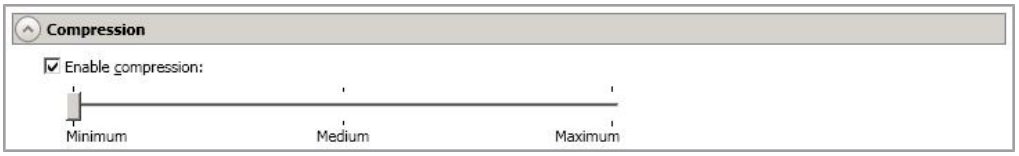

To help reduce the amount of bandwidth needed to transmit Double-Take data, compression allows you to compress data prior to transmitting it across the network. In a WAN environment this provides optimal use of your network resources. If compression is enabled, the data is compressed before it is transmitted 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

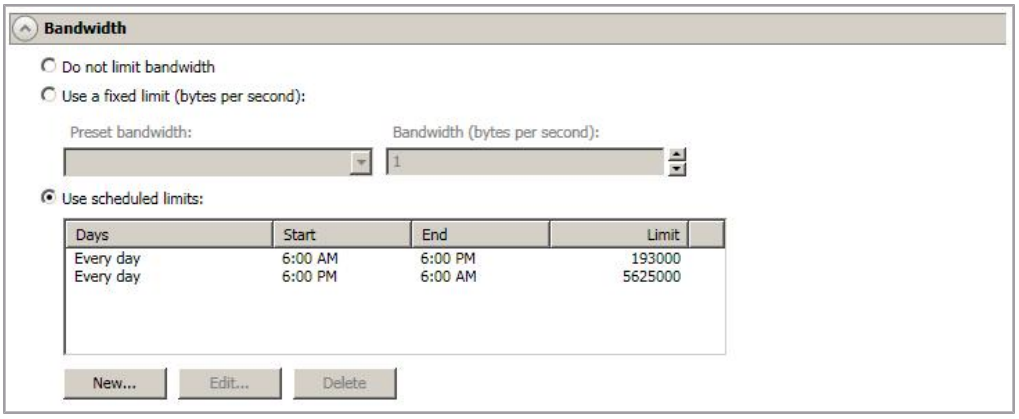

Bandwidth limitations are available to restrict the amount of network bandwidth used for Double-Take data transmissions. When a bandwidth limit is specified, Double-Take never exceeds that allotted amount. The bandwidth not in use by Double-Take is available for all other network traffic.

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

The scheduled option is not available if your source is a cluster.

- **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

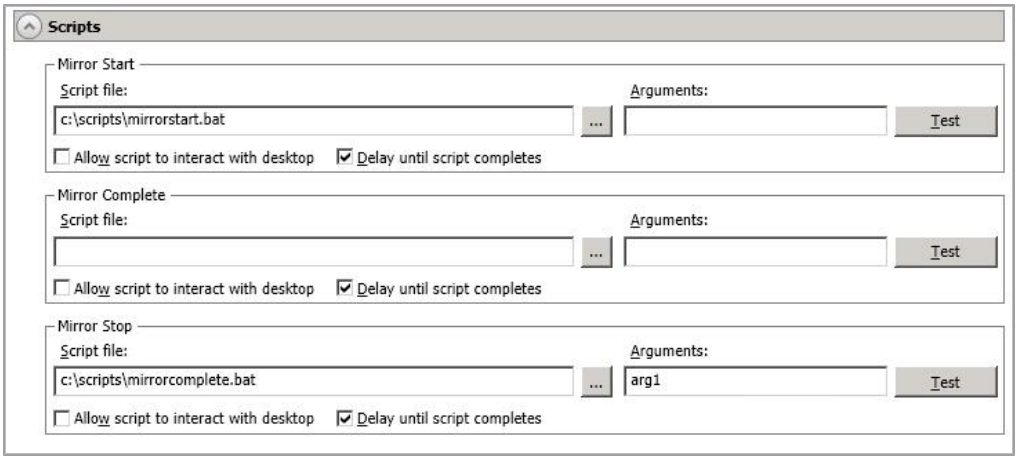

You can customize mirroring by running scripts on the target at pre-defined points in the mirroring process. Scripts may contain any valid Windows command, executable, or batch file. The scripts are processed using the same account running the Double-Take service, unless you have identified a specific account through the server's properties. See *Script credentials* on page 99. 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.
- **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.

Mirror scripts are dependent on the target and the **Target Path Mappings** specified under the **Network Route & Folder Selection** section. 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).

Test Failover Scripts

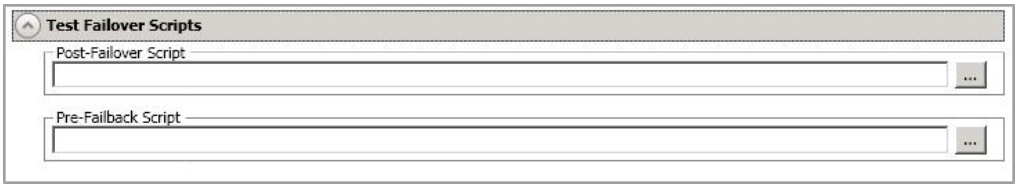

This section is not applicable to clustered environments.

When you failover, you will have three choices. You can failover to live data, failover to data from a snapshot, or perform a test failover. Any scripts you specify in this section are only used during a test failover.

- **Post-Failover Script**—This script runs on the target at the end of the failover process. Specify the full path and name of the script file.
- **Pre-Failback Script** —This script runs on the target at the beginning of the undo failover process. Specify the full path and name of the script file.

Scripts will run but will not be displayed on the screen if the Double-Take service is not set to interact with the desktop. Enable this option through the Windows Services applet.

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

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

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

Managing and controlling SQL jobs

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

The jobs displayed in the right pane depend on the server group folder selected in the left pane. Every job for each server in your console session is displayed when the **Jobs on All Servers** group is selected. If you have created and populated server groups (see *Managing servers* on page 65), then only the jobs associated with the server or target servers in that server group will be displayed in the right pane.

- See *Overview job information displayed in the top pane* on page 372
- See *Detailed job information displayed in the bottom pane* on page 374
- See *Job controls* on page 376

Overview job information displayed in the top pane

The top pane displays high-level overview information about your jobs.

Column 1 (Blank)

The first blank column indicates the state of the job.

 The job is in a healthy state.

 The job is in a warning state. This icon is also displayed on any server groups that you have created that contain a job in a warning state.

 The job is in an error state. This icon is also displayed on any server groups that you have created that contain a job in an error state.

 The job is in an unknown state.

Job

The name of the job

Source Server

The name of the source. This could be a name or IP address of a standalone server, a cluster, or a node. Cluster jobs will be associated with the cluster name and standalone jobs will be associated with a standalone server or a cluster node.

Target Server

The name of the target. This could be a name or IP address of a standalone server, a cluster, or a node. Cluster jobs will be associated with the cluster name and standalone jobs will be associated with a standalone server or a cluster node.

Job Type

Each job type has a unique job type name. This job is a SQL Server job. For a complete list of all job type names, press F1 to view the Double-Take Console online help.

Activity

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

Mirror Status

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

Replication Status

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

Transmit Mode

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

Detailed job information displayed in the bottom pane

The details displayed in the bottom pane of the **Manage Jobs** page provide additional information for the job highlighted in the top pane. If you select multiple jobs, the details for the first selected job will be displayed.

Name

The name of the job

Target data state

- **OK**—The data on the target is in a good state.
- **Mirroring**—The target is in the middle of a mirror process. The data will not be in a good state until the mirror is complete.
- **Mirror Required**—The data on the target is not in a good state because a remirror is required. This may be caused by an incomplete or stopped mirror or an operation may have been dropped on the target.
- **Restore Required**—The data on the source and target do not match because of a failover condition. Restore the data from the target back to the source. If you want to discard the changes on the target, you can remirror to resynchronize the source and target.
- **Snapshot Reverted**—The data on the source and target do not match because a snapshot has been applied on the target. Restore the data from the target back to the source. If you want to discard the changes on the target, you can remirror to resynchronize the source and target.
- **Busy**—The source is low on memory causing a delay in getting the state of the data on the target.
- **Not Loaded**—Double-Take target functionality is not loaded on the target server. This may be caused by an activation code error.
- **Unknown**—The console cannot determine the status.

Mirror remaining

The total number of mirror bytes that are remaining to be sent from the source to the target

Mirror skipped

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

Replication queue

The total number of replication bytes in the source queue

Disk queue

The amount of disk space being used to 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**.

Connected since

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

Recent activity

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

Additional information

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

Job controls

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

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

Create a New Job 

This button leaves the **Manage Jobs** page and opens the **Get Started** page.

View Job Details 

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

Delete 

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 77. You will remain on the **Manage Jobs** page after updating the server credentials. If your servers use the same credentials, make sure you also update the credentials on the **Manage Servers** page so that the Double-Take Console can authenticate to the servers in the console session. See *Managing servers* on page 65.

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. Failover monitoring will continue 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. Failover monitoring will continue while the job is stopped. 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.

Snapshots are not applicable to clustered environments.

Manage Snapshots 

Allows you to manage your snapshots by taking and deleting snapshots for the selected job. See *Managing snapshots* on page 161 for more information.

Snapshots are not applicable to clustered environments.

Failover or Cutover 

Starts the failover process. See *Failing over SQL jobs* on page 390 for the process and details of failing over a SQL job.

Failback 

Starts the failback process. See *Restoring then failing back SQL jobs* on page 392 for the process and details of failing back a SQL job.

Restore 

Starts the restoration process. See *Restoring then failing back SQL jobs* on page 392 for the process and details of restoring a SQL job.

Reverse 

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

Undo Failover 

Cancels a test failover by undoing it. Undo failover does not apply to clustered jobs. See *Failing over SQL jobs* on page 390 for details on undoing a test failover.

View Job Log 

Opens the job log. On the right-click menu, this option is called **View Logs**, and you have the option of opening the job log, source server log, or target server log. See *Viewing the log files through the Double-Take Console* on page 678 for details on all three of these logs.

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 all files**—All protected files will be mirrored from the source to the target.
  - **Mirror different files**—Only those protected files that are different based on date and time, size, or attributes will be mirrored from the source to the target.
    - **Mirror only if the file on the source is newer than the copy on the target**—Only those protected files that are newer on the source are mirrored to the target.

    

    If you are using a database application or are protecting a domain controller, do not use this option 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 that might be newer, get mirrored.

- **Use block checksum for comparisons**—For those files flagged as different, the mirroring process can perform a block checksum comparison and send only those blocks that are different.

- **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 the calculation is enabled, it is completed before the job starts mirroring, which can take a significant amount of time depending on the number of files and system performance. If your job contains a large number of files, for example, 250,000 or more, you may want to disable the calculation so that data will start being mirrored sooner. 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.
  - **Create verification report only**—This option verifies the data and generates a verification log, but it does not remirror any files that are different on the source and target. See *Verification log* on page 102 for details on the log file.
  - **Mirror files to the target server automatically**—When this option is enabled, in addition to verifying the data and generating a log, Double-Take will also mirror to the target any protected files that are different on the source.
    - **Mirror only if the file on the source is newer than the copy on the target**—Only those protected files that are newer on the source are mirrored to the target.

      

      If you are using a database application or are protecting a domain controller, do not use this option 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 that might be newer, get mirrored.

    - **Use block checksum for comparisons**—For those files flagged as different, the mirroring process can perform a block checksum comparison and send only those blocks that are different.
- **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.
- **Update Shares**—Windows share information is automatically updated on the target once an hour. This option allows you to manually update share information immediately when the option is selected. Shares are not applicable to environments where the target is a cluster.

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 65.

Type a server name

Displays only jobs that contain the text you entered. If you have created and populated server groups, then only jobs that contain the text you entered associated with the server or target servers in that server group will be displayed. See *Managing servers* on page 65.

Overflow Chevron 

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

Viewing SQL job details

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

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

Job name

The name of the job

Job type

Each job type has a unique job type name. This job is a SQL Server job. For a complete list of all job type names, press F1 to view the Double-Take Console online help.

Health

 The job is in a healthy state.

 The job is in a warning state.

 The job is in an error state.

 The job is in an unknown state.

Activity

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

Connection ID

The incremental counter used to number connections. The number is incremented when a connection is created. It is also incremented by internal actions, such as an auto-disconnect and auto-reconnect. The lowest available number (as connections are created, stopped, deleted, and so on) will always be used. The counter is reset to one each time the Double-Take service is restarted.

Transmit mode

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

Target data state

- **OK**—The data on the target is in a good state.
- **Mirroring**—The target is in the middle of a mirror process. The data will not be in a good state until the mirror is complete.
- **Mirror Required**—The data on the target is not in a good state because a remirror is required. This may be caused by an incomplete or stopped mirror or an operation may have been dropped on the target.
- **Restore Required**—The data on the source and target do not match because of a failover condition. Restore the data from the target back to the source. If you want to discard the changes on the target, you can remirror to resynchronize the source and target.
- **Snapshot Reverted**—The data on the source and target do not match because a snapshot has been applied on the target. Restore the data from the target back to the source. If you want to discard the changes on the target, you can remirror to resynchronize the source and target.
- **Busy**—The source is low on memory causing a delay in getting the state of the data on the target.
- **Not Loaded**—Double-Take target functionality is not loaded on the target server. This may be caused by an activation code error.
- **Unknown**—The console cannot determine the status.

Target route

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

Compression

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

Encryption

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

Bandwidth limit

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

Connected since

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

Additional information

Depending on the current state of your job, you may see additional information displayed to keep you informed about the progress and status of your job. If there is no

additional information, you will see (None) displayed.

Mirror status

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

Mirror percent complete

The percentage of the mirror that has been completed

Mirror remaining

The total number of mirror bytes that are remaining to be sent from the source to the target

Mirror skipped

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

Replication status

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

Replication queue

The total number of replication bytes in the source queue

Disk queue

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

Bytes sent

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

Bytes sent compressed

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

Validating a SQL job

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

1. From the **Manage Jobs** page, highlight the job and click **View Job Details** in the toolbar.
2. In the **Tasks** area on the right on the **View Job Details** page, click **Validate job properties**.
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 enable protection. Depending on the error, you may be able to click **Fix** or **Fix All** and let Double-Take correct the problem for you. For those errors that Double-Take cannot correct automatically, you will need to modify the source or target to correct the error, or you can select a different target. You must revalidate the selected servers, by clicking **Recheck**, until the validation check passes without errors.

   Validation checks for an existing job are logged to the job log on the target server. See *Log files* on page 677 for details on the various log files.
4. Once your servers have passed validation, click **Close**.

Editing a SQL job

Use these instructions to edit a SQL job.

1. From the **Manage Jobs** page, highlight the job and click **View Job Details** in the toolbar.
2. In the **Tasks** area on the right on the **View Job Details** page, click **Edit job properties**. (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 have the same options available for your SQL job as when you created the job. See *Creating a SQL job* on page 347 for details on each job option.

   

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

4. If you want to modify the workload items or replication rules for the job, click **Edit workload or replication rules**. Modify the **Workload item** you are protecting, if desired. Additionally, you can modify the specific **Replication Rules** for your job. If your job is a clustered SQL job, you will not be able to edit the workload or replication rules.

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

   You can also enter wildcard rules, however you should do so carefully. Rules are applied to files that are closest in the directory tree to them. If you have rules that include multiple folders, an exclusion rule with a wild card will need to be added for each folder that it needs applied to. For example, if you want to exclude all .log files from D:\ and your rules include D:\, D:\Dir1 , and D:\Dir2, you would need to add the exclusion rule for the root and each subfolder rule. So you will need to add exclude rules for D:*.log , D:\Dir1*.log, and D:\Dir2*.log.

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

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

   Errors are designated by a white X inside a red circle. Warnings are designated by a black exclamation point (!) inside a yellow triangle. A successful validation is designated by a white checkmark inside a green circle. You can sort the list by the icon to see errors, warnings, or successful validations together. Click on any of the validation items to see details. You must

correct any errors before you can enable protection. Depending on the error, you may be able to click **Fix** or **Fix All** and let Double-Take correct the problem for you. For those errors that Double-Take cannot correct automatically, you will need to modify the source or target to correct the error, or you can select a different target. You must revalidate the selected servers, by clicking **Recheck**, until the validation check passes without errors.

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.

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

Viewing a SQL job log

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

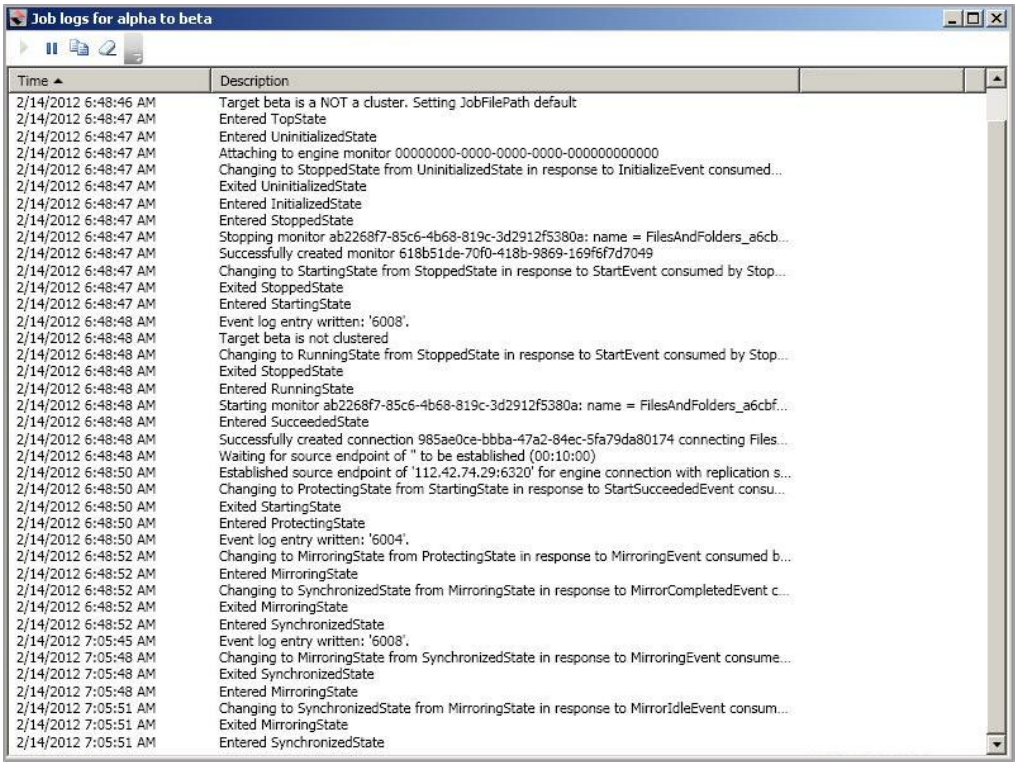

The following table identifies the controls and the table columns in the **Job logs** window.

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.

Failing over SQL jobs

When a failover condition has been met, failover will be triggered automatically if you disabled the wait for user option during your failover configuration. If the wait for user before failover option is enabled, you will be notified in the console when a failover condition has been met. At that time, you will need to trigger it manually from the console when you are ready.

1. On the **Manage Jobs** page, highlight the job that you want to failover and click **Failover or Cutover** in the toolbar.
2. Select the type of failover to perform.
   - **Failover to live data**—Select this option to initiate a full, live failover using the current data on the target. The application services will be stopped on the source (if it is online), and they will be started on the target. DNS records will be updated, if applicable. Application requests destined for the source server or its IP addresses are routed to the target.
   - **Perform test failover**—Select this option to perform a test failover using the current data on the target for standalone environments. This option is not applicable to clustered environments.This option will allow you to confirm SQL on the target is viable for failover. This process will pause the target (the entire target is not paused, just Double-Take processing), take a snapshot of the target (so that you can revert back to the pre-test state after the test is complete), and then start the application services on the target. Success and failure messages will be available in the job log. (Note the application services are still running on the source during the test.) Once the application services are running on the target, you can perform any testing desired on the target server.

     

     The following caveats apply to performing a test failover for SQL.
     - Make sure you have enough space on your target to contain the snapshot.
     - If you have scheduled Double-Take snapshots, they will continue to be created during the test, however, you should not use any of these snapshots from this time period because the target data may not be in a good state.
     - If you are using Windows 2008 R2 and Double-Take snapshots, any snapshots that you take before testing failover will be unusable for an actual failover. To work around this issue, take another snapshot immediately after the test. Once the additional, post-testing snapshot is taken, all of the previous snapshots will be usable for future failovers. This issue is fixed in the Windows 2008 R2 Service Pack 1 release. If you install that release or a later release, you will not have to worry about the extra snapshot after the test.
     - If any of your SQL data is stored on the system volume, you will not be able to perform a test failover.
     - Any scripts you specified in the **Test Failover Scripts** section when you created your job will be executed during the test failover.

   - **Failover to a snapshot**—Select this option to initiate a full, live failover without using the current data on the target. Instead, select a snapshot and the data on the target will be

reverted to that snapshot. This option will not be available if there are no snapshots on the target or if the target does not support snapshots. This option is also not applicable to clustered environments. To help you understand what snapshots are available, the **Type** indicates the kind of snapshot.

- **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. Select how you want to handle the data in the target queue. You may want to check the amount of data in queue on the target by reviewing the *Statistics* on page 688 or *Performance Monitor* on page 790.
   - **Apply data in target queues before failover or cutover**—All of the data in the target queue will be applied before failover begins. The advantage to this option is that all of the data that the target has received will be applied before failover begins. 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 failover or cutover immediately**—All of the data in the target queue will be discarded and failover will begin immediately. The advantage to this option is that failover will occur immediately. The disadvantage is that any data in the target queue will be lost.
   - **Revert to last good snapshot if target data state is bad**—If the target data is in a bad state, Double-Take will automatically revert to the last good Double-Take snapshot before failover begins. If the target data is in a good state, Double-Take will not revert the target data. Instead, Double-Take will apply the data in the target queue and then failover. The advantage to this option is that good data on the target is guaranteed to be used. The disadvantage is that if the target data state is bad, you will lose any data between the last good snapshot and the failure.
4. When you are ready to begin failover, click **Failover**.
5. If you performed a test failover, you can undo it by selecting **Undo Failover or Cutover** in the toolbar. The application services will be stopped on the target, the pre-test snapshot that was taken will be reverted, and the target will be resumed.

Restoring then failing back SQL jobs

Restoring before failing back allows your users to continue accessing their data on the failed over target, which is standing in for the source, while you perform the restoration process.

1. Resolve the problem(s) on the source that caused it to fail.
2. Bring the source onto the network.
3. On the **Manage Jobs** page, highlight the job and click **Restore**.

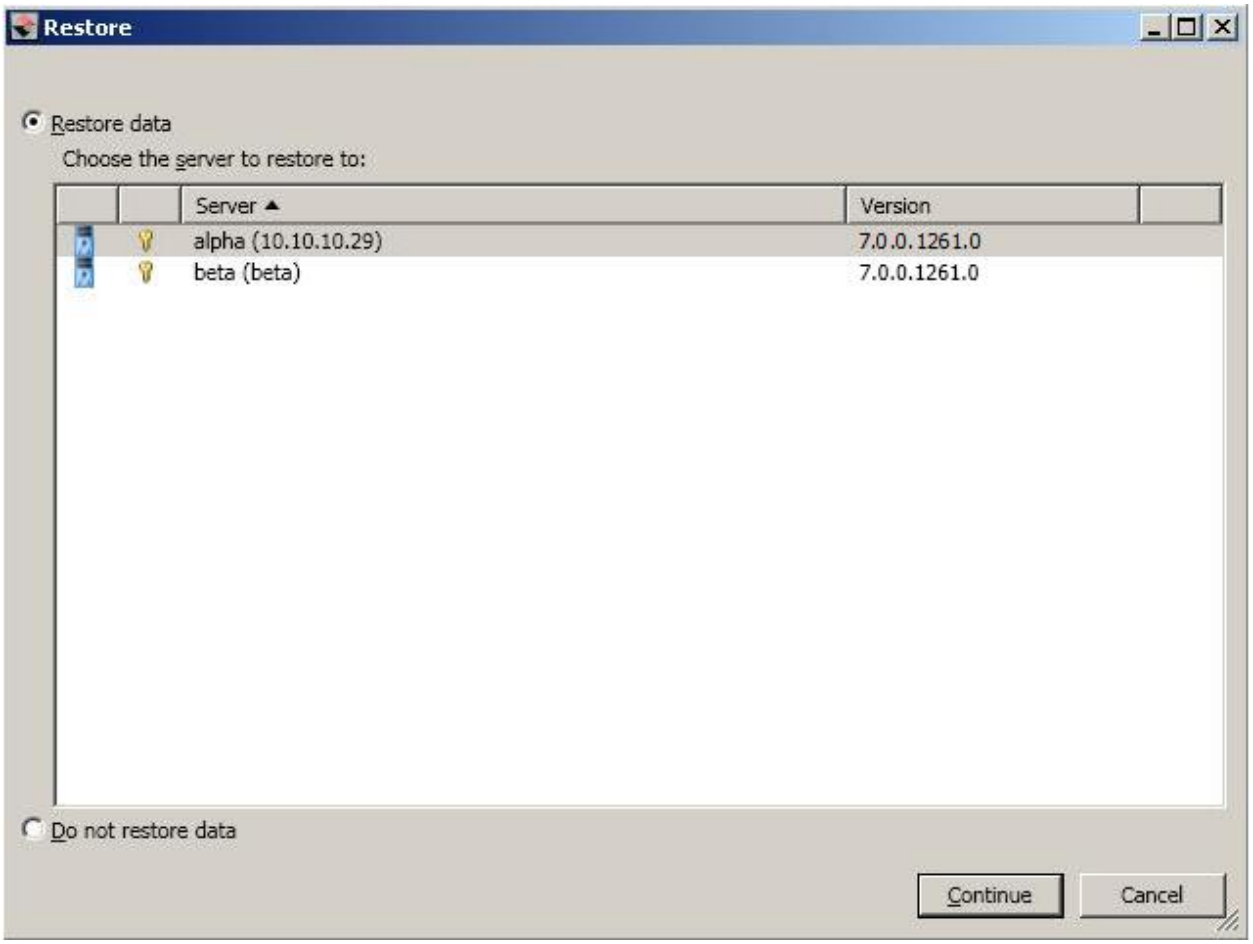

4. Confirm **Restore data** is selected, then highlight your source server in the server list.

If you do not want to restore your data, you can select **Do not restore data**. Keep in mind, that any data changes that took place on the target after failover will be lost.

---

5. Click **Continue** to start the restoration.
6. During the restoration process, the **Activity** will indicate **Restoring** and **Mirror Status** will indicate **In Progress**. When the restoration is complete, the **Mirror Status** will change to **Idle**, and the **Activity** will be **Restored**. At this time, schedule a time for failback. User downtime will begin once failback is started, so select a time that will have minimal disruption on your users.

7. On the **Manage Jobs** page, highlight the job and click **Failback**.
8. In the dialog box, highlight the job that you want to failback and click **Failback**.
9. Check that your source is fully functional and that the source data is in a good state, and then, if desired, you can enable protection again by clicking **Start**.

Chapter 11 Full server to Hyper-V protection

Create a full server to Hyper-V job when you want to protect an entire physical server or virtual machine to a Hyper-V target. There is no reverse protection for this job. You will have to use another full server job type to get back to your original hardware after failover.

- See *Full server to Hyper-V requirements* on page 395—Full server to Hyper-V protection includes specific requirements for this type of protection.
- See *Creating a full server to Hyper-V job* on page 397—This section includes step-by-step instructions for creating a full server to Hyper-V job.
- See *Managing and controlling full server to Hyper-V jobs* on page 420—You can view status information about your full server to Hyper-V jobs and learn how to control these jobs.
- See *Failing over full server to Hyper-V jobs* on page 438—Use this section when a failover condition has been met or if you want to failover manually.

Full server to Hyper-V requirements

In addition to the *Core Double-Take requirements* on page 23, use these requirements for full server to Hyper-V protection.

- **Source server**—The source server can be any physical or virtual server running any of the operating systems listed in the *Core Double-Take requirements* on page 23. However, if you are using a Windows 2003 operating system, you must have Service Pack 2 which is required for Hyper-V Integration Services. Additionally, your source cannot be a Hyper-V server.
- **Target server**—The target server can be any Windows 2008, 2008 R2, 2012, or 2012 R2 operating system from the *Core Double-Take requirements* on page 23 that has the Hyper-V role enabled. In addition, you can use Hyper-V Server 2008 R2, Server Core 2008 R2, Server Core 2012, or Server Core 2012 R2 with the Hyper-V role enabled. (Hyper-V Server 2008 and Server Core 2008 are not supported.)
- **Server Core**—In addition to the Server Core requirements above, there is a Server Core limitation. DNS updates are not supported for Server Core servers.
- **Disk types**—Virtual machines can use raw, pass-through, or differencing disks, however, they will be virtual hard disks on the replica on the target.
- **IP addressing**—IPv4 is the only supported IP version.
- **Microsoft .NET Framework**—Microsoft .NET Framework version 3.5 Service Pack 1 is required. This version is not included in the .NET version 4.0 release. Therefore, even if you have .NET version 4.0 installed, you will also need version 3.5.1. For Windows 2008 and earlier, you can install this version from the Double-Take DVD, via a web connection during the Double-Take installation, or from a copy you have obtained manually from the Microsoft web site. For Windows 2008 R2 and later, you need to enable it through Windows features.
- **Snapshots**—You can take and failover to Double-Take snapshots using a full server to Hyper-V job. See *Core Double-Take requirements* on page 23 for the specific snapshot requirements.
- **Supported configurations**—The following table identifies the supported configurations for a full server to Hyper-V job.

Configuration		Supported	Not Supported
Source to target configuration[1]	One to one, active/standby	X	
	One to one, active/active		X
	Many to one	X	
	One to many	X	
	Chained		X
	Single server		X

Configuration		Supported	Not Supported
Server configuration[2]	Standalone to standalone	X	
	Standalone to cluster		X
	Cluster to standalone	X	
	Cluster to cluster		X
	Cluster Shared Volumes (CSV) guest level	X	
	Cluster Shared Volumes (CSV) host level		X
Upgrade configuration[3]	Upgrade 5.3 P/V to Hyper-V job to 7.0 full server to Hyper-V job	X	
	Upgrade 6.0 full server to Hyper-V job to 7.0 full server to Hyper-V job	X	
Version 7.0 console[4,5]	Version 7.0 console can create job for 5.3 source and 5.3 target		X
	Version 7.0 console can create job for 6.0 source and 6.0 target	X	
	Version 7.0 console can create job for 7.0 source and 7.0 target	X	

1. See *Supported configurations* on page 16 for details on each of the source to target configurations.
2. Supported cluster configurations can be used to protect a node, shared storage, or a virtual machine, however, there is no cluster resource or cluster-awareness with this configuration.
3. When upgrading from version 5.3, you can perform a rolling upgrade where you update the target server first. After the upgrade is complete, the source will automatically reconnect to the target. At this point, the job will be an unmanaged job that you can delete or failover. No other job controls will be available. Once you upgrade you source, the job will be fully controllable.
4. Once you upgrade your console to version 7.0, existing jobs that are running version 5.3 will not appear in the console until the target of the job is upgraded to version 7.0.
5. Newer job options available in the version 7.0 console will not be functional when creating jobs for servers running version 6.0.

Creating a full server to Hyper-V job

Use these instructions to create a full server to Hyper-V job.

1. Click **Get Started** from the toolbar.
2. Select **Double-Take Availability** and click **Next**.
3. Select **Protect files and folders, an application, or an entire Windows server** and click **Next**.
4. Choose your source server. This is the physical or virtual server that you want to protect.

- **Current Servers**—This list contains the servers currently available in your console session. Servers that are not licensed for the workflow you have selected will be filtered out of the list. Select your source server from the list.
- **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 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.

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. The user must also have administrative rights for Microsoft Hyper-V.

Your source can have no more than four NICs enabled.

---

5. Click **Next** to continue.
6. Choose the type of workload that you want to protect. Under **Server Workloads**, in the **Workload types** pane, select **Full Server to Hyper-V or ESX**. 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.

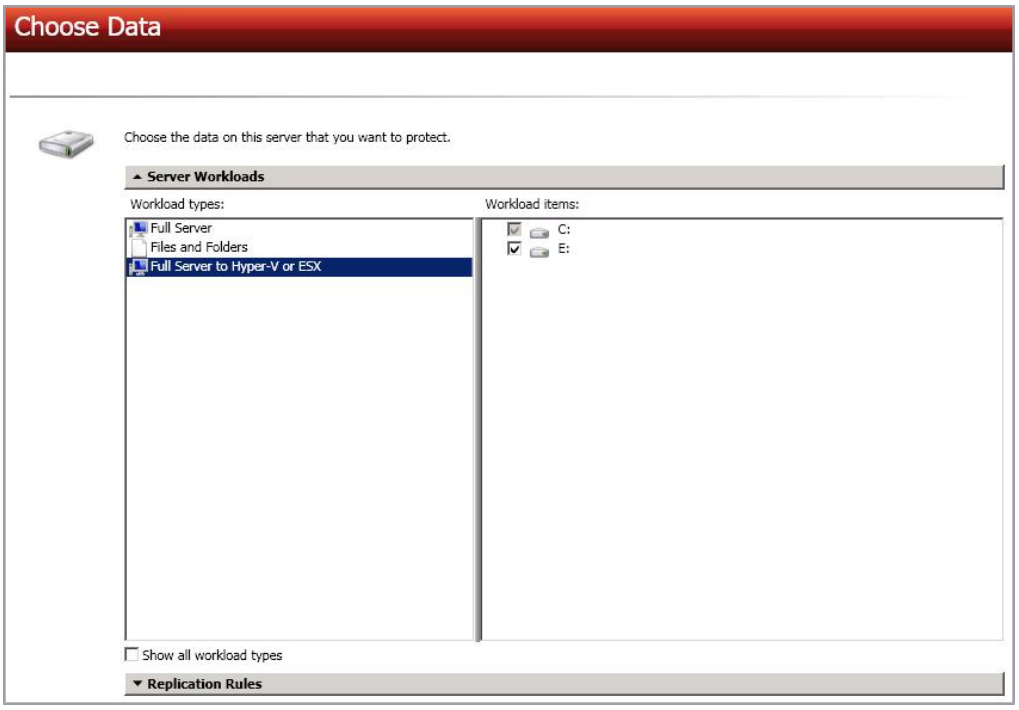

7. By default, Double-Take selects the system volume for protection. You will be unable to deselect the system volume. Select any other volumes on the source that you want to protect. If desired, click the **Replication Rules** heading and expand the volumes under **Folders**. You will see that Double-Take automatically excludes particular files that cannot be used during the protection. If desired, you can exclude other files that you do not want to protect, but be careful when excluding data. Excluded volumes, folders, and/or files may compromise the integrity of your installed applications. There are some volumes, folders, and files (identified in italics text) that you will be unable to exclude, because they are required for protection. For example, the boot files cannot be excluded because that is where the system state information is stored.

   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.

---

8. Click **Next** to continue.
9. Choose your target server. This is the Hyper-V server that will store the replica of the virtual machine from the source.

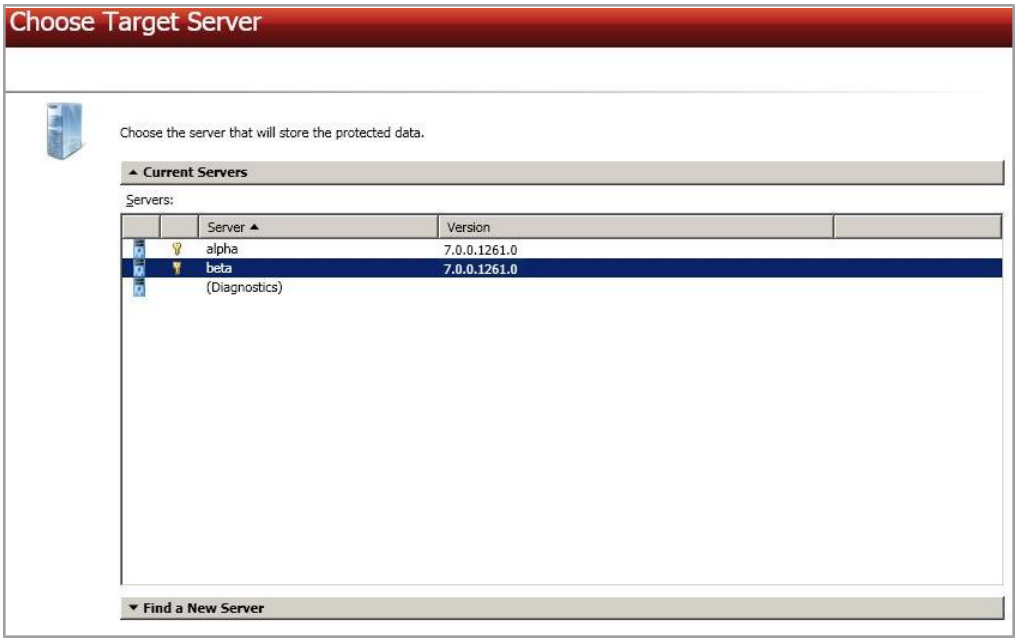

- **Current Servers**—This list contains the servers currently available in your console session. Servers that are not licensed for the workflow you have selected and those not applicable to the workload type you have selected will be filtered out of the list. Select your

target server from the list.

- **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. The user must also have administrative rights for Microsoft Hyper-V.

---

10. Click **Next** to continue.
11. You have many options available for your full server to Hyper-V 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 419.

    - *General* on page 401
    - *Replica Virtual Machine Location* on page 402
    - *Replica Virtual Machine Configuration* on page 403
    - *Replica Virtual Machine Volumes* on page 404
    - *Replica Virtual Machine Network Settings* on page 405
    - *Failover Monitor* on page 406
    - *Failover Options* on page 408
    - *Failover Identity* on page 409
    - *Mirror, Verify & Orphaned Files* on page 411
    - *Network Route* on page 415
    - *Snapshots* on page 416
    - *Compression* on page 417
    - *Bandwidth* on page 418

General

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

Replica Virtual Machine Location

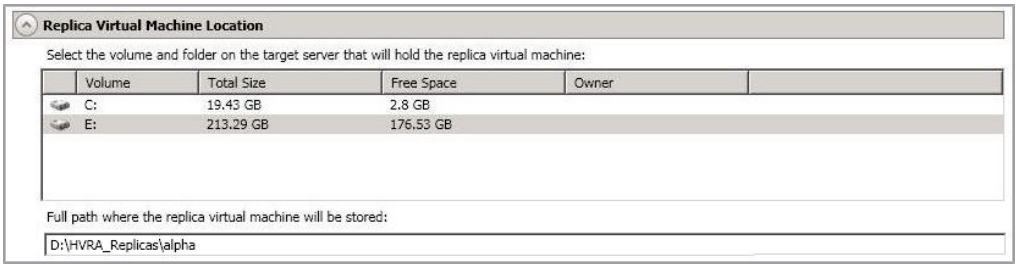

- **Select the volume and folder on the target server that will hold the replica virtual machine**—Select one of the volumes from the list to indicate the volume on the target where you want to store the new virtual server when it is created. The target volume must have enough **Free Space** to store the source data.
- **Full path where the replica virtual machine will be stored**—Specify a location on the selected **Volume** to store the replica of the source. By specifying an existing folder, you can reuse an existing virtual machine on your Hyper-V target created by a previous protection job. This can be useful for pre-staging data on a virtual machine over a LAN connection and then relocating it 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. In order to use a pre-existing virtual disk, it must be a valid virtual disk and it cannot be attached to any registered virtual machine. In a WAN environment, you may want to take advantage of re-using an existing virtual disk by using a process similar to the following.
    a. Create a protection job in a LAN environment, letting Double-Take create the virtual disk for you.
    b. Complete the mirror process locally.
    c. Delete the protection job and when prompted, select to keep the replica.
    d. From the Hyper-V Manager, delete the replica virtual machine, which will delete the virtual machine configuration but will keep the associated hard disk files.
    e. Shut down and move the Hyper-V target server to your remote site.
    f. After the Hyper-V target server is back online at the remote site, create a new protection job for the same source server. Double-Take will reuse the existing hard disk files and perform a difference mirror over the WAN to bring the virtual machine up-to-date.

Replica Virtual Machine Configuration

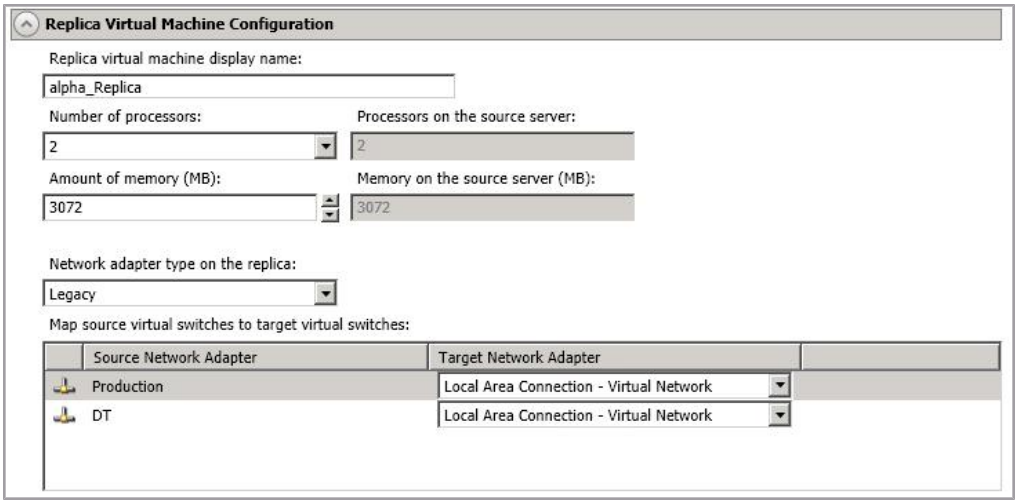

- **Replica virtual machine display name**—Specify the name of the replica virtual machine. This will be the display name of the virtual machine on the host system.
- **Number of processors**—Specify how many processors to create on the new virtual machine. The number of processors on the source is displayed to guide you in making an appropriate selection. If you select fewer processors than the source, your clients may be impacted by slower responses.
- **Amount of memory**—Specify the amount of memory, in MB, to create on the new virtual machine. The memory on the source is displayed to guide you in making an appropriate selection. If you select less memory than the source, your clients may be impacted by slower responses.
- **Network adapter type on the replica**—Depending on your operating system, you may be able to select the type of adapter, **Legacy** or **Synthetic**, to use on the replica virtual machine. This selection will apply to all adapters on the replica.
    - **Windows 2003**—This operating system only supports legacy adapters.
    - **Windows 2008 and 2008 R2**—These operating systems support both legacy and synthetic adapters.
    - **Windows 2012 and 2012 R2**—These operating systems supports both legacy and synthetic adapters, unless the target is Windows 2012 R2, in which case only the synthetic adapter is available for these operating systems.
- **Map source virtual switches to target virtual switches**—Identify how you want to handle the network mapping after failover. The **Source Network Adapter** column lists the NICs from the source. Map each one to a **Target Network Adapter**, which is a virtual network on the target. You can also choose to discard the source's NIC and IP addresses, or you can to failover the NIC and IP addresses but leave them in a not connected state.

Replica Virtual Machine Volumes

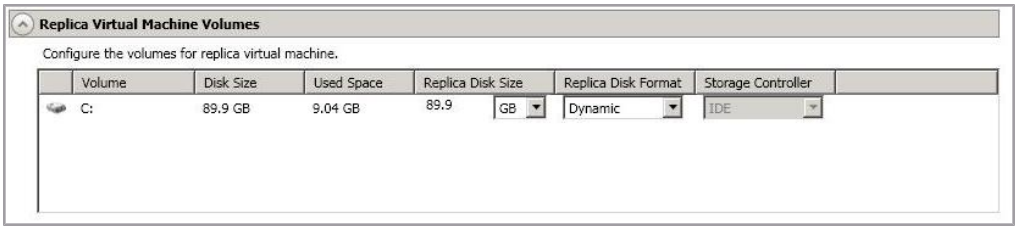

- **Replica Disk Size**—For each volume you are protecting, specify the size of the replica disk on the target. Be sure and include the value in MB or GB for the disk. The value must be at least the size of the specified **Used Space** on that volume.

In some cases, the replica virtual machine may use more virtual disk space than the size of the source volume due to differences in how the virtual disk's block size is formatted and how hard links are handled. To avoid this issue, specify the size of your replica to be at least 5 GB larger.

Snapshots are stored on the replica, so if you enable snapshots, be sure that you configure your replica virtual machine disk size large enough to maintain snapshots.

- **Replica Disk Format**—For each volume you are protecting, specify the format of the disk, **Dynamic** or **Fixed**, that will be created on the replica virtual machine. Any disk format specification will be discarded if you are reusing a disk from the **Full path where the replica virtual machine will be stored** from the **Replica Virtual Machine Location** section.
- **Storage Controller**—For each volume you are protecting, specify the type of **Storage Controller** that you want to use for each volume on the target. If your virtual machine is a Generation 2 VM (Windows 2012 or later), SCSI is the only controller option.

The system volume must be an IDE controller. In addition, up to two more volumes can be attached to an IDE controller. If you are protecting more than three volumes on the source, you will need to install the Hyper-V Integration Components to acquire a SCSI device after failover to attach these volumes to the replica virtual machine. You must be using Windows 2003 Service Pack 2 or later to use Hyper-V Integration Components.See your Microsoft documentation for more information.

Replica Virtual Machine Network Settings

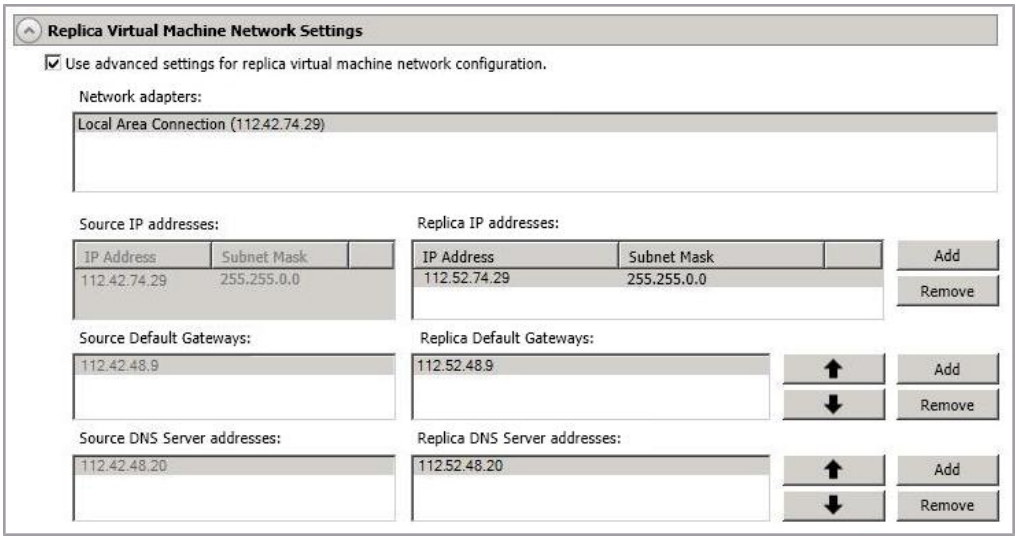

- **Use advanced settings for replica virtual machine network configuration**—Select this option to enable the replica virtual machine network setting configuration. This setting is primarily used for WAN support.
- **Network adapters**—Select a network adapter from the source and specify the **Replica IP addresses**, **Replica Default Gateways**, and **Replica DNS Server addresses** to be used after failover. If you add multiple gateways or DNS servers, you can sort them by using the arrow up and arrow down buttons. Repeat this step for each network adapter on the source.

Updates made during failover will be based on the network adapter name when protection is established. If you change that name, you will need to delete the job and re-create it so the new name will be used during failover.

If you update one of the advanced settings (IP address, gateway, or DNS server), then you must update all of them. Otherwise, the remaining items will be left blank. If you do not specify any of the advanced settings, the replica virtual machine will be assigned the same network configuration as the source.

By default, the source IP address will be included in the target IP address list as the default address. If you do not want the source IP address to be the default address on the target after failover, remove that address from the **Replica IP addresses** list.

Failover Monitor

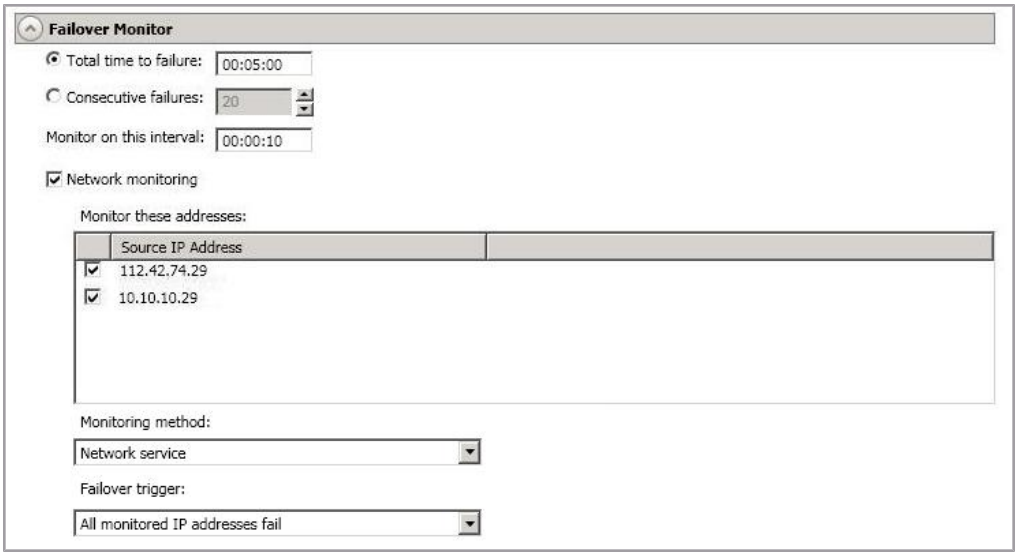

- **Total time to failure**—Specify, in hours:minutes:seconds, how long the target will keep trying to contact the source before the source is considered failed. This time is precise. If the total time has expired without a successful response from the source, this will be considered a failure.

  Consider a shorter amount of time for servers, such as a web server or order processing database, which must remain available and responsive at all times. Shorter times should be used where redundant interfaces and high-speed, reliable network links are available to prevent the false detection of failure. If the hardware does not support reliable communications, shorter times can lead to premature failover. Consider a longer amount of time for machines on slower networks or on a server that is not transaction critical. For example, failover would not be necessary in the case of a server restart.

- **Consecutive failures**—Specify how many attempts the target will make to contact the source before the source is considered failed. For example, if you have this option set to 20, and your source fails to respond to the target 20 times in a row , this will be considered a failure.
- **Monitor on this interval**—Specify, in hours:minutes:seconds, how long to wait between attempts to contact the source to confirm it is online. This means that after a response (success or failure) is received from the source, Double-Take will wait the specified interval time before contacting the source again. If you set the interval to 00:00:00, then a new check will be initiated immediately after the response is received.

  If you choose **Total time to failure**, do not specify a longer interval than failure time or your server will be considered failed during the interval period.

  If you choose **Consecutive failures**, your failure time is calculated by the length of time it takes your source to respond plus the interval time between each response, times the number of consecutive failures that can be allowed. That would be (response time + interval) * failure number. Keep in mind that timeouts from a failed check are included in the response time, so your failure time will not be precise.

- **Network monitoring**—With this option, the target will monitor the source using a network ping.
  - **Monitor these addresses**—Select each **Source IP Address** that you want the target to monitor. If you want to monitor additional addresses, enter the address and click **Add**.
  - **Monitoring method**—This option determines the type of network ping used for failover monitoring.
    - **Network service**—Source availability will be tested by an ICMP ping to confirm the route is active.
    - **Replication service**—Source availability will be tested by a UDP ping to confirm the Double-Take service is active.
    - **Network and replication services**—Source availability will be tested by both an ICMP ping to confirm the route is active and a UDP ping to confirm the Double-Take service is active. Both pings must fail in order to trigger a failover.
  - **Failover trigger**—If you are monitoring multiple IP addresses, specify when you want a failover condition to be triggered.
    - **One monitored IP address fails**—A failover condition will be triggered when any one of the monitored IP addresses fails. If each IP address is on a different subnet, you may want to trigger failover after one fails.
    - **All monitored IP addresses fail**—A failover condition will be triggered when all monitored IP addresses fail. If there are multiple, redundant paths to a server, losing one probably means an isolated network problem and you should wait for all IP addresses to fail.

Failover Options

- **Wait for user to initiate failover**—By default, the failover process will wait for you to initiate it, allowing you to control when failover occurs. When a failure occurs, the job will wait in **Failover Condition Met** for you to manually initiate the failover process. Disable this option only if you want failover to occur immediately when a failure occurs.

Failover Identity

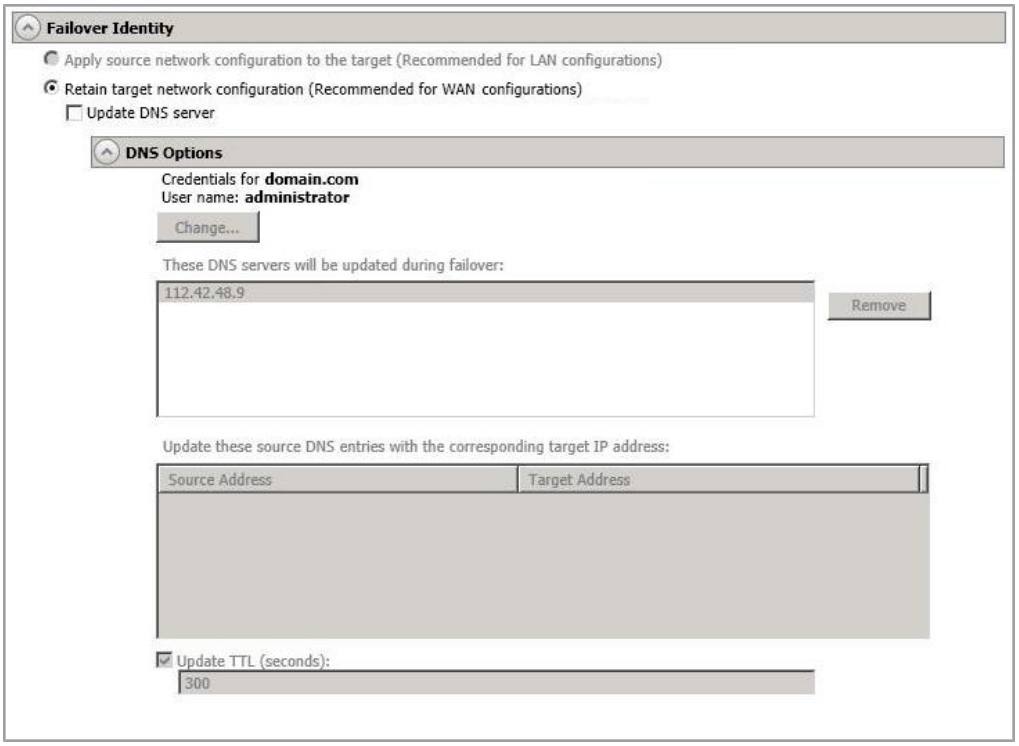

- **Retain target network configuration**—The target will retain all of its original IP addresses.
  - **Update DNS server**—Specify if you want Double-Take to update your DNS server on failover. If DNS updates are made, the DNS records will be locked during failover. Be sure and review the *Core Double-Take requirements* on page 23 for the requirements for updating DNS.

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

    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.

If you select **Retain your target network configuration** but do not enable **Update DNS server**, you will need to specify failover scripts that update your DNS server during failover, or you can update the DNS server manually after failover. This would also apply to non--Microsoft Active Directory Integrated DNS servers. You will want to keep your target network configuration but do not update DNS. In this case, you will need to specify failover scripts that update your DNS server during failover, or you can update the DNS server manually after failover.

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

Mirror, Verify & Orphaned Files

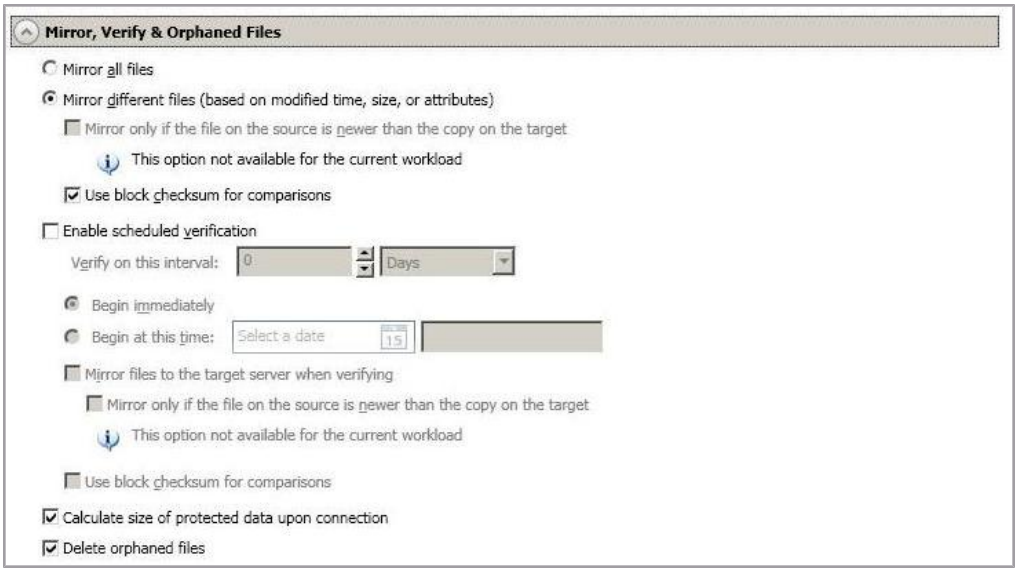

- **Mirror all files**—All protected files will be mirrored from the source to the target.
- **Mirror different files**—Only those protected files that are different based on date and time, size, or attributes will be mirrored from the source to the target.
  - **Mirror only if the file on the source is newer than the copy on the target**—This option is not available for full server to Hyper-V jobs.
  - **Use block checksum for comparisons**—For those files flagged as different, the mirroring process can perform a block checksum comparison and send only those blocks that are different.

File Differences Mirror Options Compared

The following table will help you understand how the various difference mirror options work together, including when you are using the block checksum option configured through the *Source server properties* on page 92.

An X in the table indicates that option is enabled. An X enclosed in parentheses (X) indicates that the option can be on or off without impacting the action performed during the mirror.

Not all job types have the source newer option available.

Source Server Properties	Job Properties			Action Performed
Block Checksum Option	File Differences Option	Source Newer Option	Block Checksum Option	
(X)	X			Any file that is different on the source and target based on the date, time, size, and/or attribute is transmitted to the target. The mirror sends the entire file.
(X)	X	X		Any file that is newer on the source than on the target based on date and/or time is transmitted to the target. The mirror sends the entire file.
	X		X	Any file that is different on the source and target based on date, time, size, and/or attributed is flagged as different. The mirror then performs a checksum comparison on the flagged files and only sends those blocks that are different.
X	X		X	The mirror performs a checksum comparison on all files and only sends those blocks that are different.
(X)	X	X	X	Any file that is newer on the source than on the target based on date and/or time is flagged as different. The mirror then performs a checksum comparison on the flagged files and only sends those blocks that are different.

- **Enable scheduled verification**—Verification is the process of confirming that the source replica data on the target is identical to the original 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, can automatically initiate a remirror, if configured. The remirror ensures data integrity between the source and target. When this option is enabled, Double-Take will verify the source replica data on the target and generate a verification log.

  

  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.

  - **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 at the specified date and time.
  - **Mirror files to the target server when verifying**—When this option is enabled, in addition to verifying the data and generating a log, Double-Take will also mirror to the target any protected files that are different on the source.
    - **Mirror only if the file on the source is newer than the copy on the target**—This option is not available for full server to Hyper-V jobs.
    - **Use block checksum for comparisons**—For those files flagged as different, the mirroring process can perform a block checksum comparison and send only those blocks that are different.
- **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 the calculation is enabled, it is completed before the job starts mirroring, which can take a significant amount of time depending on the number of files and system performance. If your job contains a large number of files, for example, 250,000 or more, you may want to disable the calculation so that data will start being mirrored sooner. 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 to Hyper-V jobs. The calculation time is when the system state protection processes hard links. If you disable the calculation, the hard link processing will not occur and you may have problems after failover, especially if your source is Windows 2008

- **Delete orphaned files**—An orphaned file is a file that exists in the replica data on the

target, but does not exist in the protected data on the source. This option specifies if orphaned files should be deleted on the target during a mirror, verification, or restoration.

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

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 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.

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 95 for more information.

During a mirror, orphaned file processing success messages will be logged to a separate orphaned file log. 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 100 for details on the location of that folder. The orphaned log file is overwritten during each orphaned file processing during a mirror, and the log file will be a maximum of 50 MB.

Network Route

For **Send data to the target server using this route**, Double-Take will select, by default, a target route for transmissions. If desired, specify an alternate route on the target that the data will be transmitted through. This allows you to select a different route for Double-Take traffic. For example, you can separate regular network traffic and Double-Take traffic on a machine with multiple IP addresses.

The IP address used on the source will be determined through the Windows route table.

Snapshots

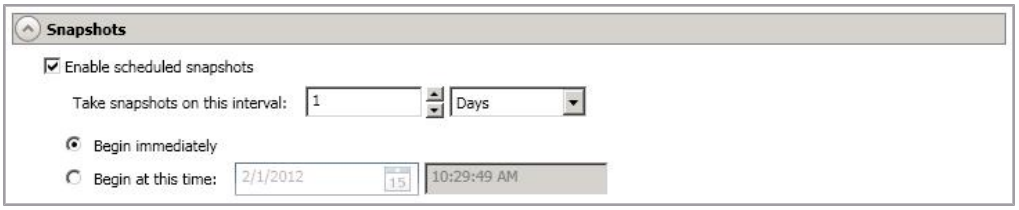

A snapshot is an image of the source replica data on the target taken at a single point in time. You can failover to a snapshot. However, you cannot access the snapshot to recover specific files or folders.

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 starting at a later date and time. Specify the date and time parameters to indicate when you want to start.

See *Managing snapshots* on page 161 for details on taking manual snapshots and deleting snapshots.

Snapshots are stored on the source replica on the target, so be sure that you configure your replica virtual machine disks large enough to maintain snapshots. Also, you may want to set the size limit on how much space snapshots can use. See your VSS documentation for more details on setting a size limit.

Compression

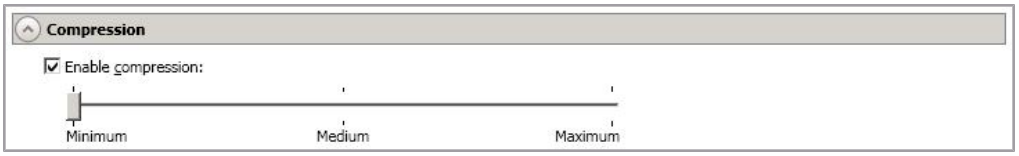

To help reduce the amount of bandwidth needed to transmit Double-Take data, compression allows you to compress data prior to transmitting it across the network. In a WAN environment this provides optimal use of your network resources. If compression is enabled, the data is compressed before it is transmitted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

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

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

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

Bandwidth

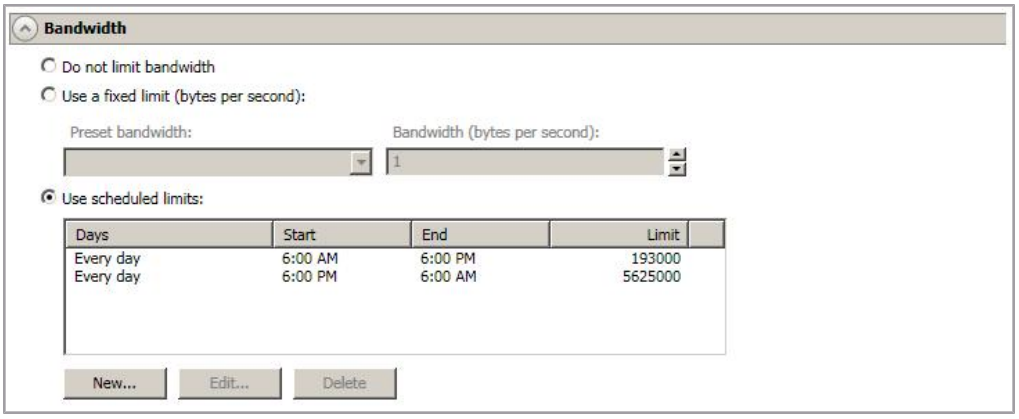

Bandwidth limitations are available to restrict the amount of network bandwidth used for Double-Take data transmissions. When a bandwidth limit is specified, Double-Take never exceeds that allotted amount. The bandwidth not in use by Double-Take is available for all other network traffic.

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

The scheduled option is not available if your source is a cluster.

- **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**.

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

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

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

Managing and controlling full server to Hyper-V jobs

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

The jobs displayed in the right pane depend on the server group folder selected in the left pane. Every job for each server in your console session is displayed when the **Jobs on All Servers** group is selected. If you have created and populated server groups (see *Managing servers* on page 65), then only the jobs associated with the server or target servers in that server group will be displayed in the right pane.

- See *Overview job information displayed in the top pane* on page 420
- See *Detailed job information displayed in the bottom pane* on page 422
- See *Job controls* on page 424

Overview job information displayed in the top pane

The top pane displays high-level overview information about your jobs.

Column 1 (Blank)

The first blank column indicates the state of the job.

 The job is in a healthy state.

 The job is in a warning state. This icon is also displayed on any server groups that you have created that contain a job in a warning state.

 The job is in an error state. This icon is also displayed on any server groups that you have created that contain a job in an error state.

 The job is in an unknown state.

Job

The name of the job

Source Server

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

Target Server

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

Job Type

Each job type has a unique job type name. This job is a Full Server to Hyper-V job. For a complete list of all job type names, press F1 to view the Double-Take Console online help.

Activity

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

Mirror Status

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

Replication Status

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

Transmit Mode

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

Detailed job information displayed in the bottom pane

The details displayed in the bottom pane of the **Manage Jobs** page provide additional information for the job highlighted in the top pane. If you select multiple jobs, the details for the first selected job will be displayed.

Name

The name of the job

Target data state

- **OK**—The data on the target is in a good state.
- **Mirroring**—The target is in the middle of a mirror process. The data will not be in a good state until the mirror is complete.
- **Mirror Required**—The data on the target is not in a good state because a remirror is required. This may be caused by an incomplete or stopped mirror or an operation may have been dropped on the target.
- **Restore Required**—The data on the source and target do not match because of a failover condition. Restore the data from the target back to the source. If you want to discard the changes on the target, you can remirror to resynchronize the source and target.
- **Snapshot Reverted**—The data on the source and target do not match because a snapshot has been applied on the target. Restore the data from the target back to the source. If you want to discard the changes on the target, you can remirror to resynchronize the source and target.
- **Busy**—The source is low on memory causing a delay in getting the state of the data on the target.
- **Not Loaded**—Double-Take target functionality is not loaded on the target server. This may be caused by an activation code error.
- **Unknown**—The console cannot determine the status.

Mirror remaining

The total number of mirror bytes that are remaining to be sent from the source to the target

Mirror skipped

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

Replication queue

The total number of replication bytes in the source queue

Disk queue

The amount of disk space being used to 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**.

Connected since

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

Recent activity

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

Additional information

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

Job controls

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

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

Create a New Job 

This button leaves the **Manage Jobs** page and opens the **Get Started** page.

View Job Details 

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

Delete 

Stops (if running) and deletes the selected jobs.

If you no longer want to protect the source and no longer need the replica of the source on the target, select to delete the associated replica virtual machine. Selecting this option will remove the job and completely delete the replica virtual machine on the target.

If you no longer want to mirror and replicate data from the source to the target but still want to keep the replica of the source on the target, select to keep the associated replica virtual machine. You may want to use this option to relocate the virtual hard disks and create a new job between the original source and the new location. Selecting this option, will preserve the source replica on the target.

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 77. You will remain on the **Manage Jobs** page after updating the server credentials. If your servers use the same credentials, make sure you also update the credentials on the **Manage Servers** page so that the Double-Take Console can authenticate to the servers in the console session. See *Managing servers* on page 65.

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. Failover monitoring will continue 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. Failover monitoring will continue while the job is stopped.

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 161 for more information.

Failover or Cutover 

Starts the failover process. See *Failing over full server to Hyper-V jobs* on page 438 for the process and details of failing over a full server to Hyper-V job.

Failback 

Starts the failback process. Failback does not apply to full server to Hyper-V jobs.

Restore 

Starts the restoration process. Restore does not apply to full server to Hyper-V jobs.

Reverse 

Reverses protection. Reverse protection does not apply to full server to Hyper-V jobs.

Undo Failover 

Cancels a test failover by undoing it. This resets the servers and the job back to their original state. See *Failing over full server to Hyper-V jobs* on page 438 for the process and details of undoing a failed over full server to Hyper-V job.

View Job Log 

Opens the job log. On the right-click menu, this option is called **View Logs**, and you have the option of opening the job log, source server log, or target server log. See *Viewing the log files through the Double-Take Console* on page 678 for details on all three of these logs.

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 all files**—All protected files will be mirrored from the source to the target.
  - **Mirror different files**—Only those protected files that are different based on date and time, size, or attributes will be mirrored from the source to the target.
    - **Mirror only if the file on the source is newer than the copy on the target**—This option is available for full server to Hyper-V jobs, but

ideally it should not be used.

- **Use block checksum for comparisons**—For those files flagged as different, the mirroring process can perform a block checksum comparison and send only those blocks that are different.

- **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 the calculation is enabled, it is completed before the job starts mirroring, which can take a significant amount of time depending on the number of files and system performance. If your job contains a large number of files, for example, 250,000 or more, you may want to disable the calculation so that data will start being mirrored sooner. 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 to Hyper-V jobs. The calculation time is when the system state protection processes hard links. If you disable the calculation, the hard link processing will not occur and you may have problems after failover, especially if your source is Windows 2008.

- **Verify**—Even if you have scheduled the verification process, you can run it manually any time a mirror is not in progress.
  - **Create verification report only**—This option verifies the data and generates a verification log, but it does not remirror any files that are different on the source and target. See *Verification log* on page 102 for details on the log file.
  - **Mirror files to the target server automatically**—When this option is enabled, in addition to verifying the data and generating a log, Double-Take will also mirror to the target any protected files that are different on the source.
    - **Mirror only if the file on the source is newer than the copy on the target**—This option is available for full server to Hyper-V jobs, but ideally it should not be used.
    - **Use block checksum for comparisons**—For those files flagged as different, the mirroring process can perform a block checksum comparison and send only those blocks that are different.
- **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.

- **Update Shares**—Shares are not applicable because they are automatically included with the system state that is being protected with the entire server.

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 65.

Type a server name

Displays only jobs that contain the text you entered. If you have created and populated server groups, then only jobs that contain the text you entered associated with the server or target servers in that server group will be displayed. See *Managing servers* on page 65.

Overflow Chevron 

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

Viewing full server to Hyper-V job details

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

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

Job name

The name of the job

Job type

Each job type has a unique job type name. This job is a Full Server to Hyper-V job. For a complete list of all job type names, press F1 to view the Double-Take Console online help.

Health

 The job is in a healthy state.

 The job is in a warning state.

 The job is in an error state.

 The job is in an unknown state.

Activity

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

Connection ID

The incremental counter used to number connections. The number is incremented when a connection is created. It is also incremented by internal actions, such as an auto-disconnect and auto-reconnect. The lowest available number (as connections are created, stopped, deleted, and so on) will always be used. The counter is reset to one each time the Double-Take service is restarted.

Transmit mode

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

Target data state

- **OK**—The data on the target is in a good state.
- **Mirroring**—The target is in the middle of a mirror process. The data will not be in a good state until the mirror is complete.
- **Mirror Required**—The data on the target is not in a good state because a remirror is required. This may be caused by an incomplete or stopped mirror or an operation may have been dropped on the target.
- **Restore Required**—The data on the source and target do not match because of a failover condition. Restore the data from the target back to the source. If you want to discard the changes on the target, you can remirror to resynchronize the source and target.
- **Snapshot Reverted**—The data on the source and target do not match because a snapshot has been applied on the target. Restore the data from the target back to the source. If you want to discard the changes on the target, you can remirror to resynchronize the source and target.
- **Busy**—The source is low on memory causing a delay in getting the state of the data on the target.
- **Not Loaded**—Double-Take target functionality is not loaded on the target server. This may be caused by an activation code error.
- **Unknown**—The console cannot determine the status.

Target route

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

Compression

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

Encryption

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

Bandwidth limit

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

Connected since

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

Additional information

Depending on the current state of your job, you may see additional information displayed to keep you informed about the progress and status of your job. If there is no

additional information, you will see (None) displayed.

Mirror status

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

Mirror percent complete

The percentage of the mirror that has been completed

Mirror remaining

The total number of mirror bytes that are remaining to be sent from the source to the target

Mirror skipped

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

Replication status

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

Replication queue

The total number of replication bytes in the source queue

Disk queue

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

Bytes sent

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

Bytes sent compressed

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

Validating a full server to Hyper-V job

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

1. From the **Manage Jobs** page, highlight the job and click **View Job Details** in the toolbar.
2. In the **Tasks** area on the right on the **View Job Details** page, click **Validate job properties**.
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 enable protection. Depending on the error, you may be able to click **Fix** or **Fix All** and let Double-Take correct the problem for you. For those errors that Double-Take cannot correct automatically, you will need to modify the source or target to correct the error, or you can select a different target. You must revalidate the selected servers, by clicking **Recheck**, until the validation check passes without errors.

   Validation checks for an existing job are logged to the job log on the target server. See *Log files* on page 677 for details on the various log files.
4. Once your servers have passed validation, click **Close**.

Editing a full server to Hyper-V job

Use these instructions to edit a full server to Hyper-V job.

1. From the **Manage Jobs** page, highlight the job and click **View Job Details** in the toolbar.
2. In the **Tasks** area on the right on the **View Job Details** page, click **Edit job properties**. (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 full server to Hyper-V job as when you created the job, but you will not be able to edit all of them. If desired, edit those options that are configurable for an existing job. See *Creating a full server to Hyper-V job* on page 397 for details on each job option.

   

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

4. If you want to modify the workload items or replication rules for the job, click **Edit workload or replication rules**. Modify the **Workload item** you are protecting, if desired. Additionally, you can modify the specific **Replication Rules** for your job.

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

   You can also enter wildcard rules, however you should do so carefully. Rules are applied to files that are closest in the directory tree to them. If you have rules that include multiple folders, an exclusion rule with a wild card will need to be added for each folder that it needs applied to. For example, if you want to exclude all .log files from D:\ and your rules include D:\, D:\Dir1 , and D:\Dir2, you would need to add the exclusion rule for the root and each subfolder rule. So you will need to add exclude rules for D:*.log , D:\Dir1*.log, and D:\Dir2*.log.

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

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

   Errors are designated by a white X inside a red circle. Warnings are designated by a black exclamation point (!) inside a yellow triangle. A successful validation is designated by a white checkmark inside a green circle. You can sort the list by the icon to see errors, warnings, or

successful validations together. Click on any of the validation items to see details. You must correct any errors before you can enable protection. Depending on the error, you may be able to click **Fix** or **Fix All** and let Double-Take correct the problem for you. For those errors that Double-Take cannot correct automatically, you will need to modify the source or target to correct the error, or you can select a different target. You must revalidate the selected servers, by clicking **Recheck**, until the validation check passes without errors.

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.

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

Viewing a full server to Hyper-V job log

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

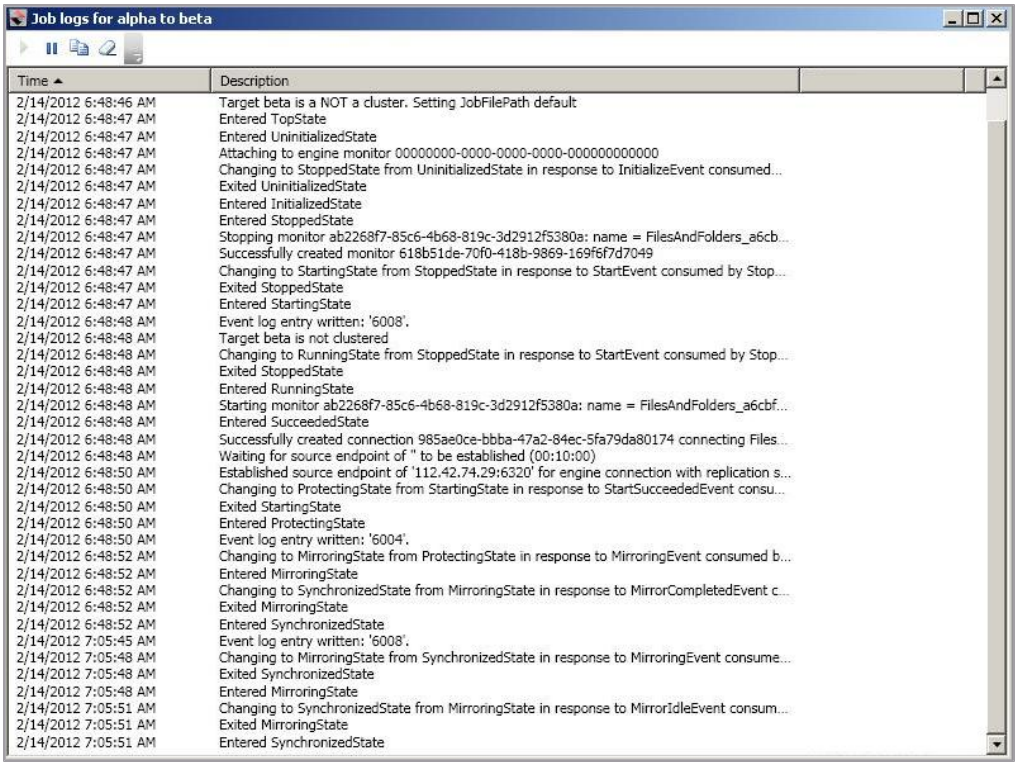

The following table identifies the controls and the table columns in the **Job logs** window.

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.

Failing over full server to Hyper-V jobs

When a failover condition has been met, failover will be triggered automatically if you disabled the wait for user option during your failover configuration. If the wait for user before failover option is enabled, you will be notified in the console when a failover condition has been met. At that time, you will need to trigger it manually from the console when you are ready.

1. On the **Manage Jobs** page, highlight the job that you want to failover and click **Failover or Cutover** in the toolbar.
2. Select the type of failover to perform.
    - **Failover to live data**—Select this option to initiate a full, live failover using the current data on the target. This option will shutdown the source machine (if it is online), stop the protection job, and start the replica virtual machine on the target with full network connectivity.
    - **Perform test failover**—Select this option to perform a test failover using the current data on the target. This option will leave the source machine online, stop the protection job, and start the replica virtual machine on the target without network connectivity.
    - **Failover to a snapshot**—Select this option to initiate a full, live failover without using the current data on the target. Instead, select a snapshot and the data on the target will be reverted to that snapshot. This option will not be available if there are no snapshots on the target or if the target does not support snapshots. This option is also not applicable to clustered environments. To help you understand what snapshots are available, the **Type** indicates the kind of snapshot.
        - **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. Select how you want to handle the data in the target queue. You may want to check the amount of data in queue on the target by reviewing the *Statistics* on page 688 or *Performance Monitor* on page 790.
    - **Apply data in target queues before failover or cutover**—All of the data in the target queue will be applied before failover begins. The advantage to this option is that all of the data that the target has received will be applied before failover begins. 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 failover or cutover immediately**—All of the data in the target queue will be discarded and failover will begin immediately. The advantage to this option is that failover will occur immediately. The disadvantage is that any data in the target queue will be lost.
    - **Revert to last good snapshot if target data state is bad**—If the target data is in a bad state, Double-Take will automatically revert to the last good Double-Take snapshot before failover begins. If the target data is in a good state, Double-Take will not revert the target data. Instead, Double-Take will apply the data in the target queue and then failover. The advantage to this option is that good data on the target is guaranteed to be used. The

disadvantage is that if the target data state is bad, you will lose any data between the last good snapshot and the failure.

4. When you are ready to begin failover, click **Failover**.

Because the Windows product activation is dependent on hardware, you may need to reactivate your Windows registration after failover. In most cases when you are using Windows 2003, you can follow the on-screen prompts to complete the reactivation. However, when you are using Windows 2008, the reactivation depends on several factors including service pack level, Windows edition, and your licensing type. If a Windows 2008 target comes online after failover 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.

Depending on your replica configuration, you may have to reboot your replica after failover. You will be prompted to reboot if it is necessary.

After failover, if you attempt to use a full server job to revert back to your original configuration, you will need to perform a few additional tasks before creating the full server job. Contact technical support if you need assistance with these steps.

1. On either the source or target, stop the Double-Take and Double-Take Management Service services.

2. Remove the GUID value from HKEY_LOCAL_MACHINE\SOFTWARE\NSI Software\Double-Take\CurrentVersion\Communication\UniqueId. Do not delete the UniqueId key. Only delete the GUI value within the key.
3. Restart the the Double-Take and Double-Take Management Service services.
4. Remove and then add your servers back into the Double-Take Console.
5. Install a different license on the original source and complete a host transfer if necessary.

If your source is running on Windows Server 2008 and the target replica has one or more SCSI drives, then after failover the CD/DVD ROM will not be allocated. If the CD/DVD ROM is required, you will need to edit the virtual machine settings to add a CD/DVD ROM after failover. By not allocating a CD/DVD ROM under these specific conditions, drive letter consistency will be guaranteed.

---

5. If you performed a test failover, you can undo it by selecting **Undo Failover or Cutover** in the toolbar. The replica virtual machine on the target will be shut down and the protection job will be restarted performing a file differences mirror.
6. There is no reverse or failback once you have failed over. If you need to go back to your original hardware, delete the job and re-create a new one if your original source was a virtual server. If your original source was a physical server, you will need to use a full server job.

Chapter 12 Full server to ESX protection

Create a full server to ESX job when you want to protect an entire physical server or virtual machine to an ESX target. There is no reverse protection for this job. You will have to use another full server job type to get back to your original hardware after failover.

- See *Full server to ESX requirements* on page 442—Full server to ESX protection includes specific requirements for this type of protection.
- See *Creating a full server to ESX job* on page 445—This section includes step-by-step instructions for creating a full server to ESX job.
- See *Managing and controlling full server to ESX jobs* on page 471—You can view status information about your full server to ESX jobs and learn how to control these jobs.
- See *Failing over full server to ESX jobs* on page 489—Use this section when a failover condition has been met or if you want to failover manually.

Full server to ESX requirements

In addition to the *Core Double-Take requirements* on page 23, use these requirements for full server to ESX protection.

- **Source server**—The source server can be any physical server running any of the operating systems listed in the *Core Double-Take requirements* on page 23. The source server must have Double-Take installed and licensed on it.
- **ESX server**—The ESX server that will host your target can be any of the following operating systems.
  - ESX 4.0.x or 4.1 Standard, Advanced, Enterprise, or Enterprise Plus
  - ESXi 4.0.x or 4.1 Standard, Advanced, Enterprise, or Enterprise Plus
  - ESXi 5.0 Standard, Enterprise, or Enterprise Plus
  - ESXi 5.1 Essentials, Essentials Plus, Standard, Enterprise, or Enterprise Plus
  - ESXi 5.5 Essentials, Essentials Plus, Standard, Enterprise, or Enterprise Plus

  

  If you are using the Standard edition of ESX 4.0 or ESXi 4.0, you must have update 1 or later.

  If your source is Windows 2008 R2, your ESX server must have ESX 4.0 update 1 or later.

  If your source is Windows 2012 or 2012 R2, your ESX server must have ESXi 5.0 update 1 or later.

- **vCenter**—vCenter is not required, but if you are using it, then you must use version 4.1 or later.
- **vMotion**—Host vMotion is only supported if you are using vCenter. Storage vMotion is not supported.
- **Virtual recovery appliance**—The ESX server must have an existing virtual machine, known as a virtual recovery appliance, that meets the following requirements. (When you establish protection, the virtual recovery appliance will create a new virtual server, mount disks, format disks, and so on. If failover occurs, the new virtual machine is detached from the virtual recovery appliance and powered on. Once the new virtual machine is online, it will have the identity, data, and system state of the source. Since the virtual recovery appliance maintains its own identity, it can be reused for additional failovers.)
  - **Operating system**—The virtual recovery appliance can be any of the operating systems listed in the *Core Double-Take requirements* on page 23.
  - **Operating system version**—The virtual recovery appliance must have the same or newer operating system than the source (not including service pack level).
  - **Operating system installation location**—Because VMware boots from the first bootable volume that is discovered, the operating system must be installed to SCSI device 0, Slot 0 on the virtual recovery appliance.
  - **Double-Take**—The virtual recovery appliance must have Double-Take installed and licensed on it.

- **Microsoft .NET Framework**—Microsoft .NET Framework version 3.5 Service Pack 1 is required on the virtual recovery appliance. This version is not included in the .NET version 4.0 release. Therefore, even if you have .NET version 4.0 installed, you will also need version 3.5.1. If you are using Windows 2008 or earlier, you can install this version from the Double-Take DVD, via a web connection during the Double-Take installation, or from a copy you have obtained manually from the Microsoft web site.If you are using Windows 2008 R2 or later, you can enable it through Windows features.
- **Snapshots**—Do not take snapshots of the virtual recovery appliance, because they will interfere with proper failover.

- **Disk controller**—VMware Paravirtual SCSI Controllers are not supported.
- **IP addressing**—IPv4 is the only supported IP version.
- **Domain controllers**—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.
- **Snapshots**—You can take and failover to Double-Take snapshots using a full server to ESX job. See *Core Double-Take requirements* on page 23 for the specific snapshot requirements.
- **Supported configurations**—The following table identifies the supported configurations for a full server to ESX job.

Configuration		Supported	Not Supported
Source to target configuration[1]	One to one, active/standby	X	
	One to one, active/active		X
	Many to one	X	
	One to many	X	
	Chained		X
	Single server		X
Server configuration[2]	Standalone to standalone	X	
	Standalone to cluster		X
	Cluster to standalone	X	
	Cluster to cluster		X
	Cluster Shared Volumes (CSV) guest level	X	
	Cluster Shared Volumes (CSV) host level		X

Configuration		Supported	Not Supported
Upgrade configuration[3]	Upgrade 5.3 P/V to ESX job to 7.0 full server to ESX job	X	
	Upgrade 6.0 full server to ESX job to 7.0 full server to ESX job	X	
Version 7.0 console[4,5]	Version 7.0 console can create job for 5.3 source and 5.3 target		X
	Version 7.0 console can create job for 6.0 source and 6.0 target	X	
	Version 7.0 console can create job for 7.0 source and 7.0 target	X	

1. See *Supported configurations* on page 16 for details on each of the source to target configurations.
2. Supported cluster configurations can be used to protect a node, shared storage, or a virtual machine, however, there is no cluster resource or cluster-awareness with this configuration.
3. When upgrading from version 5.3, you can perform a rolling upgrade where you update the target server first. After the upgrade is complete, the source will automatically reconnect to the target. At this point, the job will be an unmanaged job that you can delete or failover. No other job controls will be available. Once you upgrade you source, the job will be fully controllable.
4. Once you upgrade your console to version 7.0, existing jobs that are running version 5.3 will not appear in the console until the target of the job is upgraded to version 7.0.
5. Newer job options available in the version 7.0 console will not be functional when creating jobs for servers running version 6.0.

Creating a full server to ESX job

Use these instructions to create a full server to ESX job.

1. Click **Get Started** from the toolbar.
2. Select **Double-Take Availability** and click **Next**.
3. Select **Protect files and folders, an application, or an entire Windows server** and click **Next**.
4. Choose your source server. This is the physical or virtual server that you want to protect.

- **Current Servers**—This list contains the servers currently available in your console session. Servers that are not licensed for the workflow you have selected will be filtered out of the list. Select your source server from the list.
- **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 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.

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.

Your source can have no more than ten NICs enabled.

5. Click **Next** to continue.
6. Choose the type of workload that you want to protect. Under **Server Workloads**, in the **Workload types** pane, select **Full Server to Hyper-V or ESX**. 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.

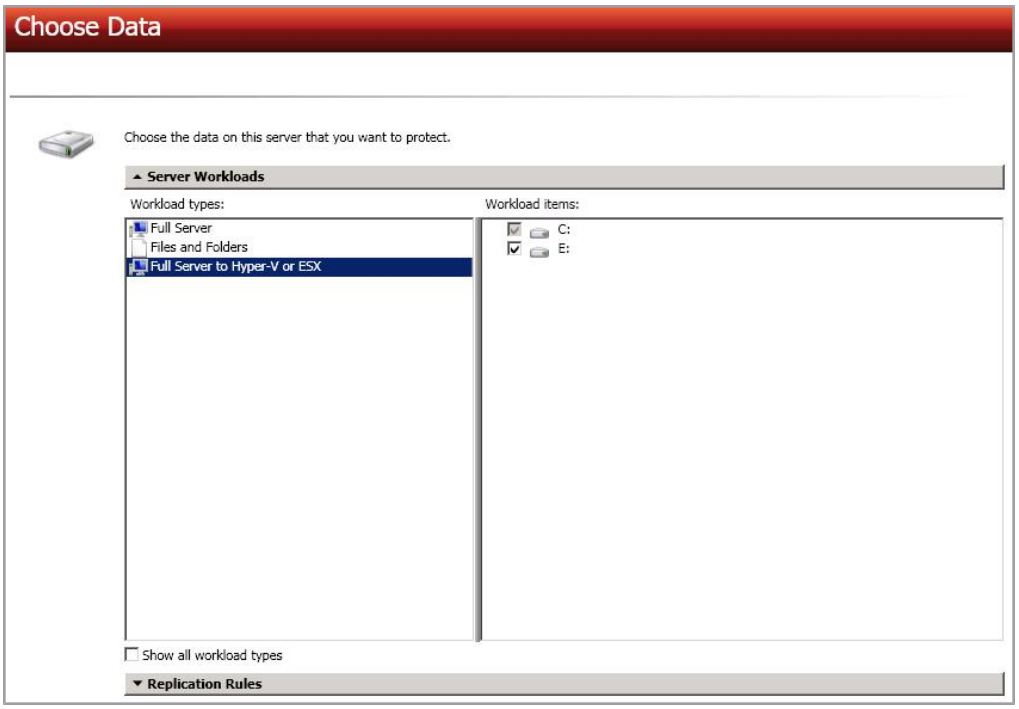

7. By default, Double-Take selects the system volume for protection. You will be unable to deselect the system volume. Select any other volumes on the source that you want to protect. If desired, click the **Replication Rules** heading and expand the volumes under **Folders**. You will see that Double-Take automatically excludes particular files that cannot be used during the protection. If desired, you can exclude other files that you do not want to protect, but be careful when excluding data. Excluded volumes, folders, and/or files may compromise the integrity of your installed applications. There are some volumes, folders, and files (identified in italics text) that you will be unable to exclude, because they are required for protection. For example, the boot files cannot be excluded because that is where the system state information is stored.

   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.

---

8. Click **Next** to continue.
9. Choose your target server. This is the virtual recovery appliance on your ESX server. See *Full server to ESX requirements* on page 442 for details on the virtual recovery appliance.

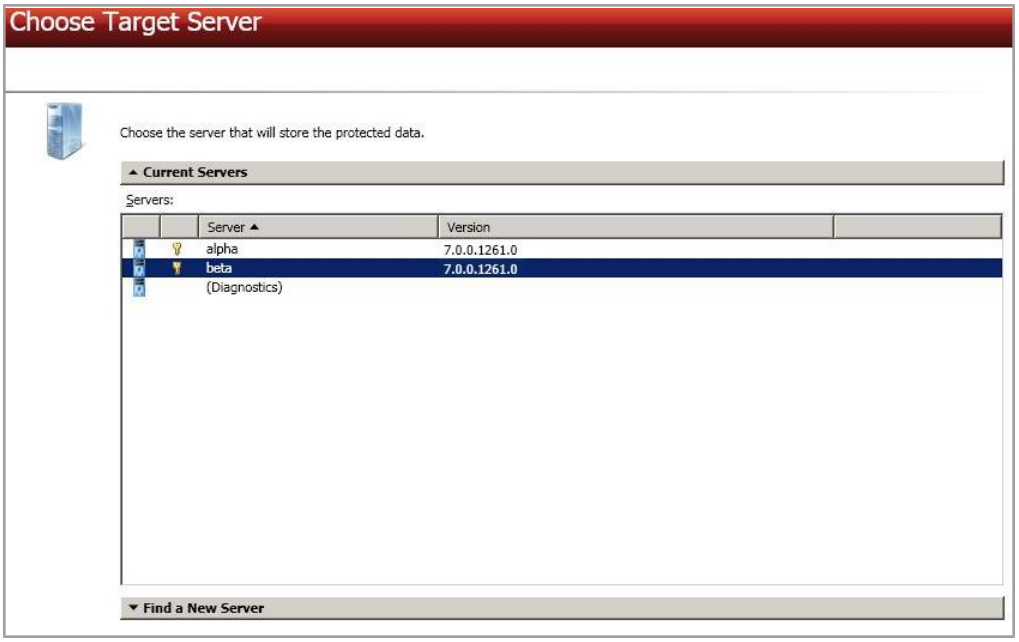

- **Current Servers**—This list contains the servers currently available in your console session. Servers that are not licensed for the workflow you have selected and those not applicable to the workload type you have selected will be filtered out of the list. Select your target server from the list.
- **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.

10. Click **Next** to continue.
11. Choose the ESX server where your target virtual recovery appliance is located. This is also the server where your replica virtual machine will be located.

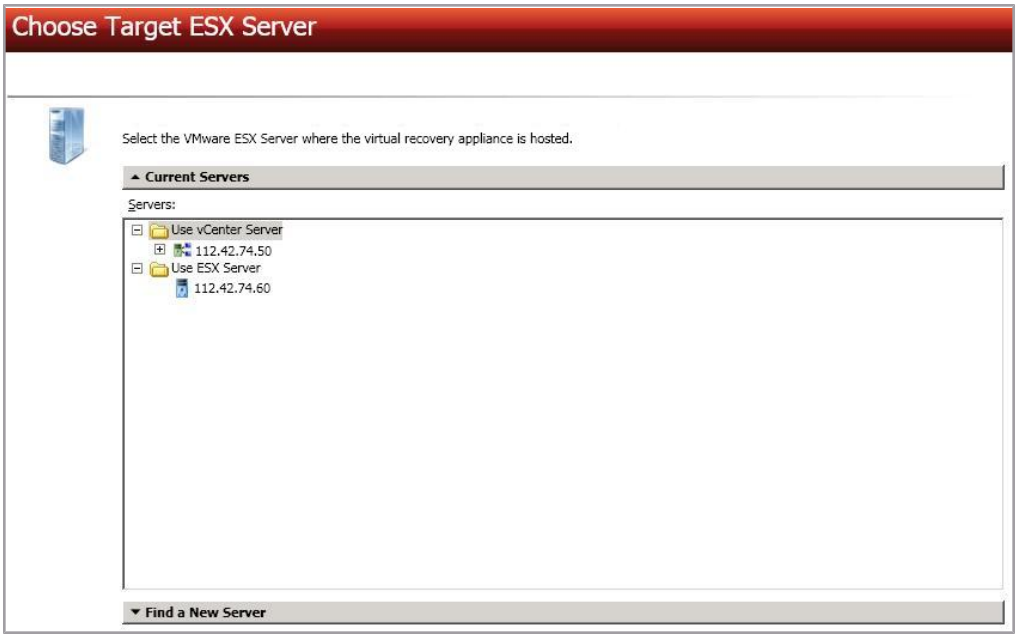

- **Current Servers**—This list contains the vCenter and ESX servers currently available in your console session. Select your ESX server from the list.
- **Find a New Server**—If the server you need is not in the **Current Servers** list, click the **Find a New Server** heading.
  - **vCenter Server**—Select your vCenter server from the list. If your vCenter server is not in the list, click **Add VirtualCenter Server**, specify the server and valid credentials, and click **Add**. If you are not using vCenter, select **None**.
  - **ESX Server**—Specify the name or IP address of the ESX server.
  - **User name**—This field will only be available if you are not using vCenter. In this case, specify the root user or another user that has the administrator role on the specified ESX server.
  - **Password**—Specify the password associated with the **User name** you entered.
  - **Domain**—If you are working in a domain environment, specify the Domain.

12. Click **Next** to continue.

13. You have many options available for your full server to ESX 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 470.

    - *General* on page 450
    - *Replica Virtual Machine Location* on page 451
    - *Replica Virtual Machine Configuration* on page 452
    - *Replica Virtual Machine Volumes* on page 453
    - *Replica Virtual Machine Network Settings* on page 456
    - *Failover Monitor* on page 457
    - *Failover Options* on page 459
    - *Failover Identity* on page 460
    - *Mirror, Verify & Orphaned Files* on page 462
    - *Network Route* on page 466
    - *Snapshots* on page 467
    - *Compression* on page 468
    - *Bandwidth* on page 469

General

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

Replica Virtual Machine Location

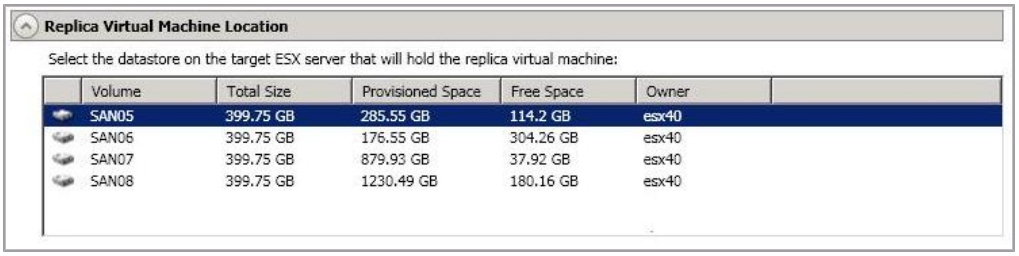

Select one of the volumes from the list to indicate the volume on the target where you want to store the configuration files for the new virtual server when it is created. The target volume must have enough **Free Space**. You can select the location of the .vmdk files under **Replica Virtual Machine Volumes**.

Replica Virtual Machine Configuration

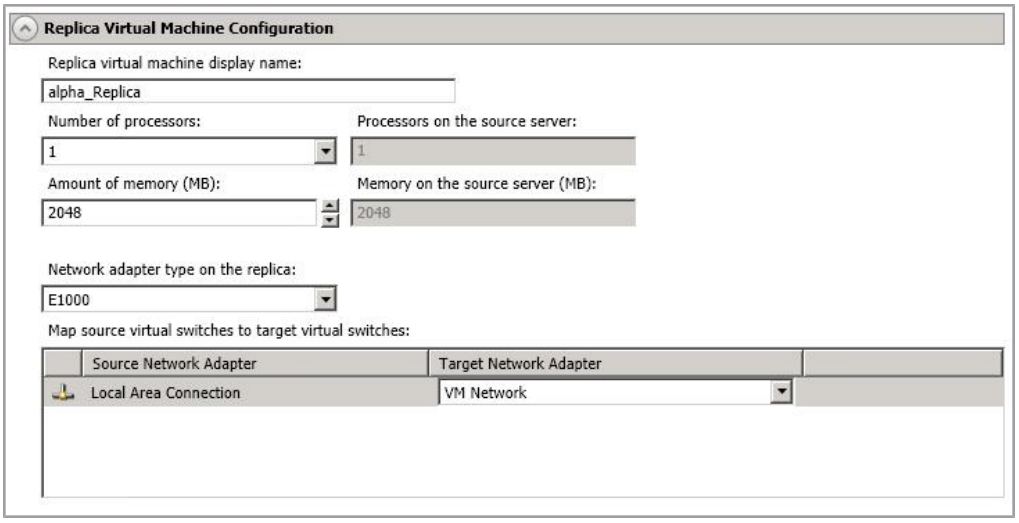

- **Replica virtual machine display name**—Specify the name of the replica virtual machine. This will be the display name of the virtual machine on the host system.
- **Number of processors**—Specify how many processors to create on the new virtual machine. The number of processors on the source is displayed to guide you in making an appropriate selection. If you select fewer processors than the source, your clients may be impacted by slower responses.
- **Amount of memory**—Specify the amount of memory, in MB, to create on the new virtual machine. The memory on the source is displayed to guide you in making an appropriate selection. If you select less memory than the source, your clients may be impacted by slower responses.
- **Network adapter type on the replica**—If you are using 2008 or later and your target appliance has VMware Tools installed, you can select the type of adapter, **E1000** or **VmxNet3**, to use on the replica virtual machine. This selection will apply to all adapters on the replica.
- **Map source virtual switches to target virtual switches**—Identify how you want to handle the network mapping after failover. The **Source Network Adapter** column lists the NICs from the source. Map each one to a **Target Network Adapter**, which is a virtual network on the target. You can also choose to discard the source's NIC and IP addresses.

Replica Virtual Machine Volumes

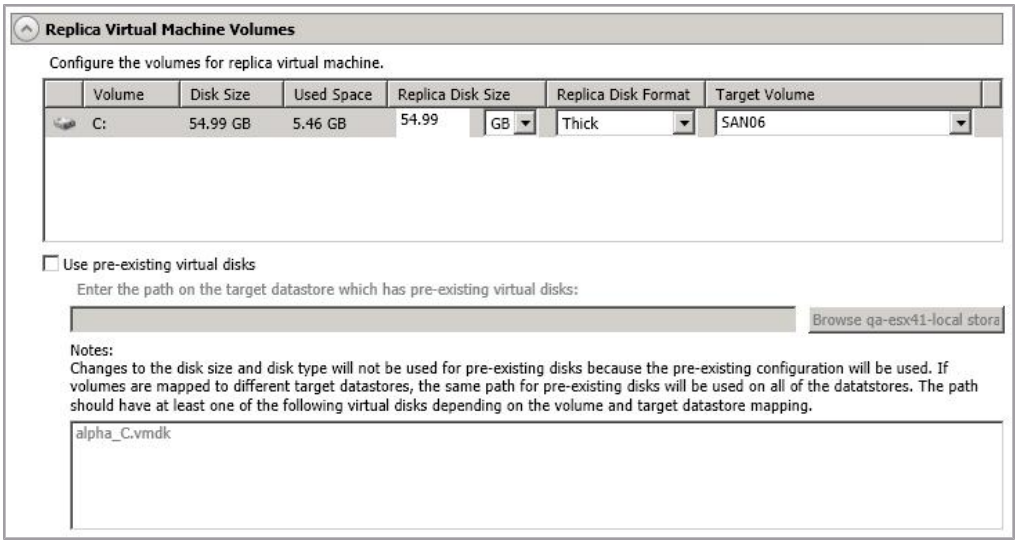

- **Replica Disk Size**—For each volume you are protecting, specify the size of the replica disk on the target. Be sure and include the value in MB or GB for the disk. The value must be at least the size of the specified **Used Space** on that volume. Any disk size specification will be discarded if you select **Use pre-existing virtual disks**, because the pre-existing configuration will be used.

  In some cases, the replica virtual machine may use more virtual disk space than the size of the source volume due to differences in how the virtual disk's block size is formatted and how hard links are handled. To avoid this issue, specify the size of your replica to be at least 5 GB larger.

  Snapshots are stored on the replica, so if you enable snapshots, be sure that you configure your replica virtual machine disk size large enough to maintain snapshots.

- **Replica Disk Format**—For each volume you are protecting, specify the format of the disk that will be created. Any disk format specification will be discarded if you select **Use pre-existing virtual disks**, because the pre-existing configuration will be used.
  - **Flat**—This disk format allocates the full amount of the disk space immediately, but does not initialize the disk space to zero until it is needed. This disk format is only available on ESX 5.
  - **Thick**—This disk format allocates the full amount of the disk space immediately, initializing all of the allocated disk space to zero.
  - **Thin**—This disk format does not allocate the disk space until it is needed.
- **Target Volume**—For each volume you are protecting, specify the volume on the target where you want to store the .vmdk files for the new replica virtual machine. If you select **Use pre-existing virtual disks**, all of the pre-existing disks must be on the same volume. You can specify the location of the virtual machine configuration files under **Replica Virtual Machine Location**.

- **Use pre-existing virtual disks**—You can reuse an existing virtual disk on your target, rather than having Double-Take create a virtual disk for you. This can be useful for pre-staging data on a virtual machine over a LAN connection and then relocating it 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. In order to use a pre-existing virtual disk, it must be a valid virtual disk. It cannot be attached to any other virtual machine, and the virtual disk size and format cannot be changed.

  Each pre-existing disk must be placed in a temporary location on the appropriate datastore, and each temporary location must be the same name. For example, a valid configuration would be datastore1/prestage, datastore2/prestage, and datastore3/prestage, but an invalid configuration would be datastore1/prestage1, datastore2/prestage2, and datastore3/prestage3. Double-Take will skip the virtual disk creation steps when using a pre-existing disk, and will instead move your existing virtual disks to the appropriate VMware location on that datastore for the virtual machine. Therefore, it is important that you place your pre-existing virtual disks in the temporary locations so this move process will be handled correctly. Specify this temporary location for **Enter the path on the target datastore which has pre-existing virtual disks**.

  In order for Double-Take to find the pre-existing disk, the virtual disk file names must be formatted using the convention SourceServer_DriveLetter. For example, if your source server is Alpha and you are protecting drives C and D, Double-Take will look for the file names Alpha_C.vmdk and Alpha_D.vmdk. If you are using IP addresses, substitute the IP address for the server name. For example, if the IP address for server Alpha is 172.31.10.25 then Double-Take will look for the file names 172.31.10.25_C.vmdk and 172.31.10.25_D.vmdk.

  If you originally created a virtual disk and specified the source server by its IP address, the pre-existing virtual disk file name cannot use the server name. However, you can rename that file and its associated -flat.vmdk file to use the IP address. The reverse is also true. If you originally specified the source server by its name, the pre-existing virtual disk file name cannot use the server's IP address. However, you can rename the file and its associated -flat.vmdk to use the source name. For example, if you originally created a virtual disk and specified the source by its IP address, you need to rename the file source_name_drive.vmdk to source_IPaddress_drive.vmdk. You also need to rename the file source_name_drive-flat.vmdk to source_IPaddress_drive-flat.vmdk. The reverse (change source_IPaddress to source_name for both files) is also true. Additionally, you will need to edit the .vmdk file manually because it contains the name of the -flat.vmdk file. Modify the reference to the -flat.vmdk file to the new name you have specified using any standard text editor.

  In a WAN environment, you may want to take advantage of the **Use pre-existing virtual disks** feature by using a process similar to the following.

  a. Create a protection job in a LAN environment, letting Double-Take create the virtual disk for you.
  b. Complete the mirror process locally.
  c. Delete the protection job and when prompted, select to keep the replica.
  d. Remove the virtual machine from the ESX inventory, which will delete the virtual machine configuration but will keep the associated.vmdk files.
  e. Shut down and move the ESX target server to your remote site.

f. After the ESX target server is back online at the remote site, move the .vmdk files to a temporary location.

g. Create a new protection job for the same source server and select to **Use pre-existing virtual disks**, specifying the temporary location of your .vmdk files. Double-Take will reuse the existing .vmdk files (automatically moving the files to the correct location) and perform a difference mirror over the WAN to bring the virtual machine up-to-date.

Replica Virtual Machine Network Settings

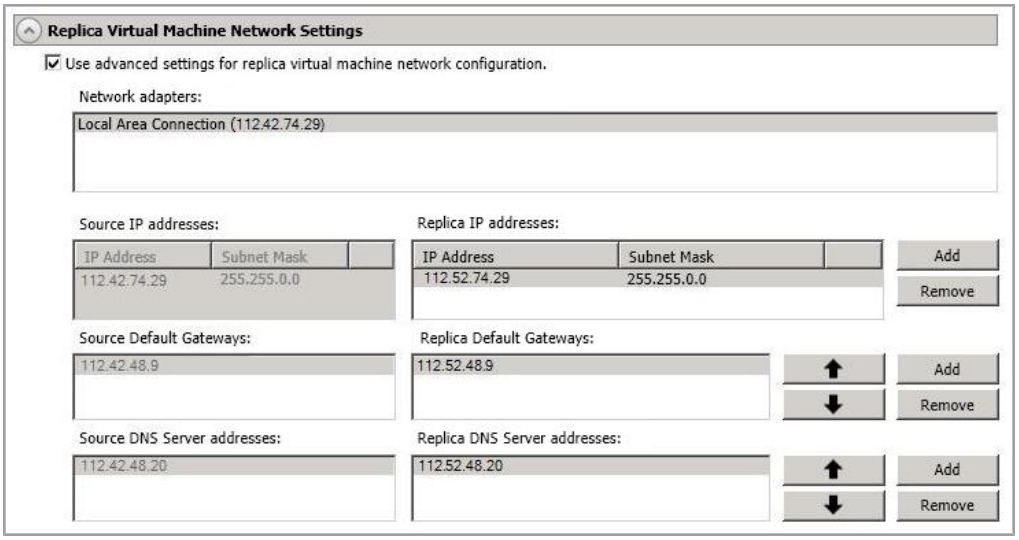

- **Use advanced settings for replica virtual machine network configuration**—Select this option to enable the replica virtual machine network setting configuration. This setting is primarily used for WAN support.
- **Network adapters**—Select a network adapter from the source and specify the **Replica IP addresses**, **Replica Default Gateways**, and **Replica DNS Server addresses** to be used after failover. If you add multiple gateways or DNS servers, you can sort them by using the arrow up and arrow down buttons. Repeat this step for each network adapter on the source.

Updates made during failover will be based on the network adapter name when protection is established. If you change that name, you will need to delete the job and re-create it so the new name will be used during failover.

If you update one of the advanced settings (IP address, gateway, or DNS server), then you must update all of them. Otherwise, the remaining items will be left blank. If you do not specify any of the advanced settings, the replica virtual machine will be assigned the same network configuration as the source.

By default, the source IP address will be included in the target IP address list as the default address. If you do not want the source IP address to be the default address on the target after failover, remove that address from the **Replica IP addresses** list.

Failover Monitor

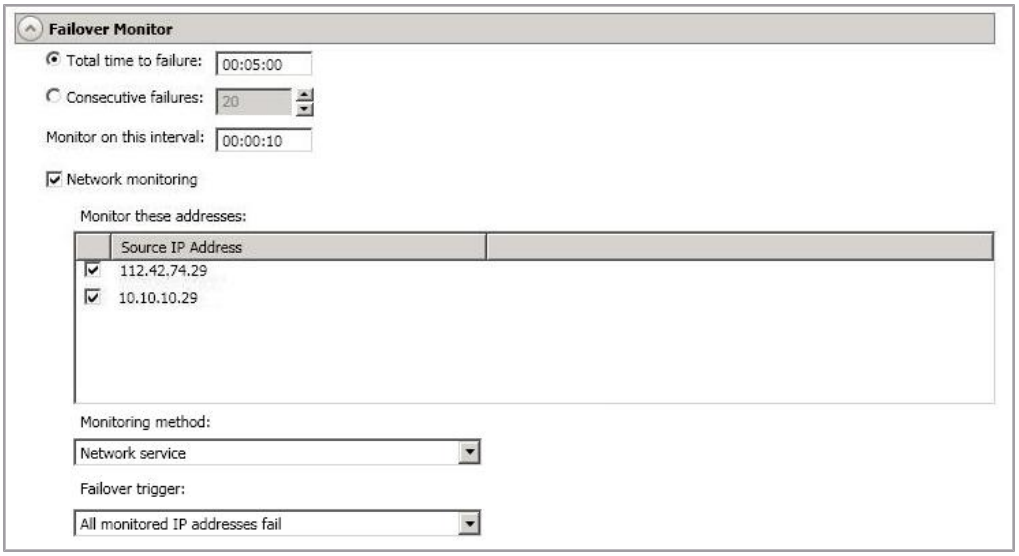

- **Total time to failure**—Specify, in hours:minutes:seconds, how long the target will keep trying to contact the source before the source is considered failed. This time is precise. If the total time has expired without a successful response from the source, this will be considered a failure.

  Consider a shorter amount of time for servers, such as a web server or order processing database, which must remain available and responsive at all times. Shorter times should be used where redundant interfaces and high-speed, reliable network links are available to prevent the false detection of failure. If the hardware does not support reliable communications, shorter times can lead to premature failover. Consider a longer amount of time for machines on slower networks or on a server that is not transaction critical. For example, failover would not be necessary in the case of a server restart.

- **Consecutive failures**—Specify how many attempts the target will make to contact the source before the source is considered failed. For example, if you have this option set to 20, and your source fails to respond to the target 20 times in a row , this will be considered a failure.
- **Monitor on this interval**—Specify, in hours:minutes:seconds, how long to wait between attempts to contact the source to confirm it is online. This means that after a response (success or failure) is received from the source, Double-Take will wait the specified interval time before contacting the source again. If you set the interval to 00:00:00, then a new check will be initiated immediately after the response is received.

  If you choose **Total time to failure**, do not specify a longer interval than failure time or your server will be considered failed during the interval period.

  If you choose **Consecutive failures**, your failure time is calculated by the length of time it takes your source to respond plus the interval time between each response, times the number of consecutive failures that can be allowed. That would be (response time + interval) * failure number. Keep in mind that timeouts from a failed check are included in the response time, so your failure time will not be precise.

- **Network monitoring**—With this option, the target will monitor the source using a network ping.
  - **Monitor these addresses**—Select each **Source IP Address** that you want the target to monitor. If you want to monitor additional addresses, enter the address and click **Add**.
  - **Monitoring method**—This option determines the type of network ping used for failover monitoring.
    - **Network service**—Source availability will be tested by an ICMP ping to confirm the route is active.
    - **Replication service**—Source availability will be tested by a UDP ping to confirm the Double-Take service is active.
    - **Network and replication services**—Source availability will be tested by both an ICMP ping to confirm the route is active and a UDP ping to confirm the Double-Take service is active. Both pings must fail in order to trigger a failover.
  - **Failover trigger**—If you are monitoring multiple IP addresses, specify when you want a failover condition to be triggered.
    - **One monitored IP address fails**—A failover condition will be triggered when any one of the monitored IP addresses fails. If each IP address is on a different subnet, you may want to trigger failover after one fails.
    - **All monitored IP addresses fail**—A failover condition will be triggered when all monitored IP addresses fail. If there are multiple, redundant paths to a server, losing one probably means an isolated network problem and you should wait for all IP addresses to fail.

Failover Options

- **Wait for user to initiate failover**—By default, the failover process will wait for you to initiate it, allowing you to control when failover occurs. When a failure occurs, the job will wait in **Failover Condition Met** for you to manually initiate the failover process. Disable this option only if you want failover to occur immediately when a failure occurs.

Failover Identity

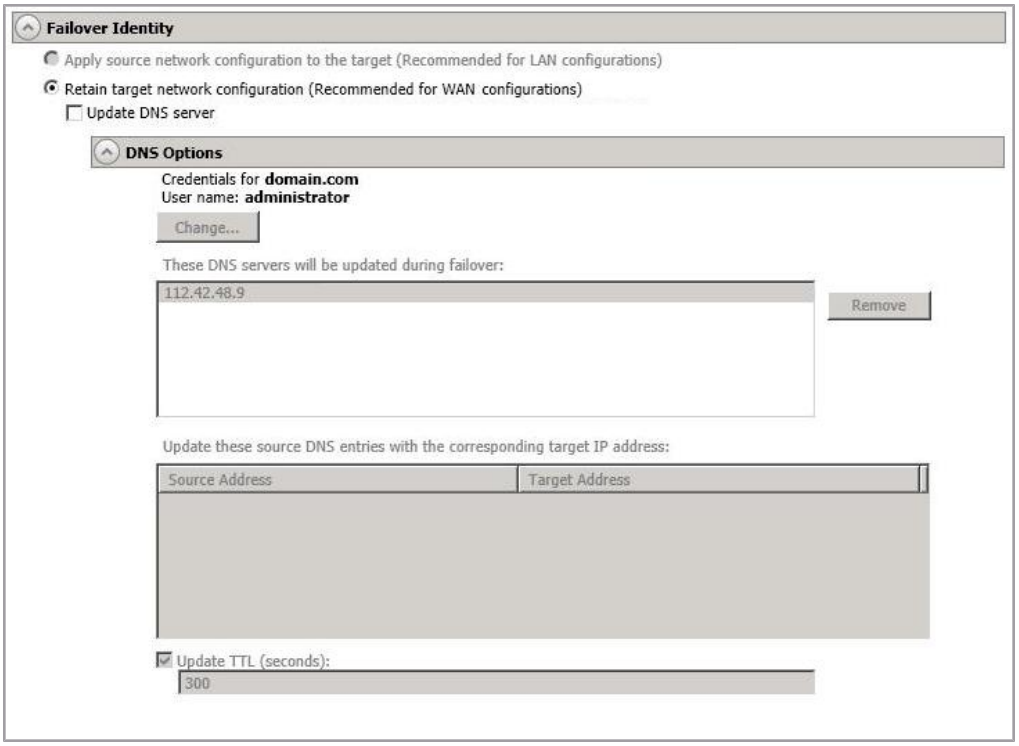

- **Retain target network configuration**—The target will retain all of its original IP addresses.
  - **Update DNS server**—Specify if you want Double-Take to update your DNS server on failover. If DNS updates are made, the DNS records will be locked during failover. Be sure and review the *Core Double-Take requirements* on page 23 for the requirements for updating DNS.

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

    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.

If you select **Retain your target network configuration** but do not enable **Update DNS server**, you will need to specify failover scripts that update your DNS server during failover, or you can update the DNS server manually after failover. This would also apply to non--Microsoft Active Directory Integrated DNS servers. You will want to keep your target network configuration but do not update DNS. In this case, you will need to specify failover scripts that update your DNS server during failover, or you can update the DNS server manually after failover.

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

Mirror, Verify & Orphaned Files

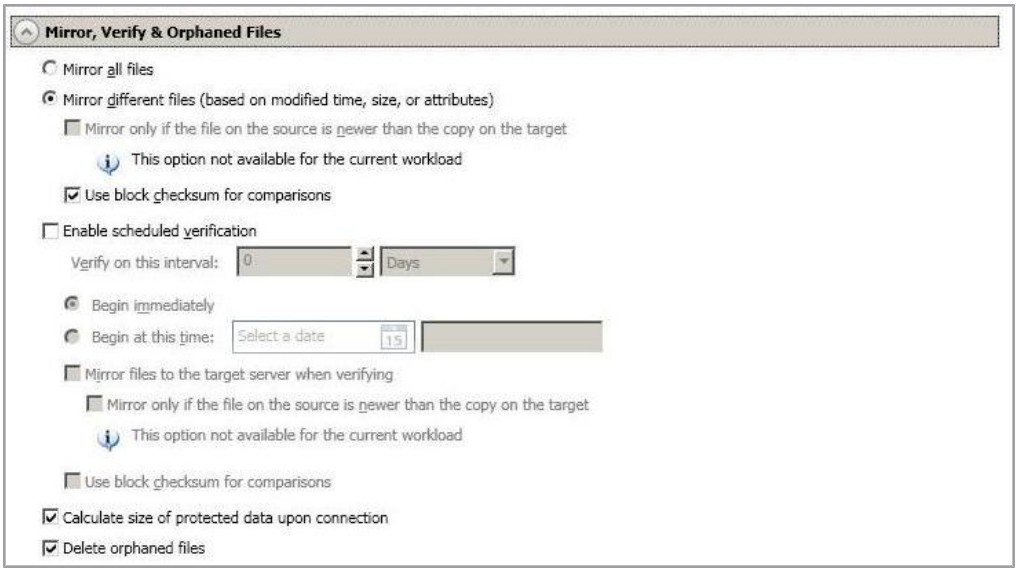

- **Mirror all files**—All protected files will be mirrored from the source to the target.
- **Mirror different files**—Only those protected files that are different based on date and time, size, or attributes will be mirrored from the source to the target.
  - **Mirror only if the file on the source is newer than the copy on the target**—This option is not available for full server to ESX jobs.
  - **Use block checksum for comparisons**—For those files flagged as different, the mirroring process can perform a block checksum comparison and send only those blocks that are different.

File Differences Mirror Options Compared

The following table will help you understand how the various difference mirror options work together, including when you are using the block checksum option configured through the *Source server properties* on page 92.

An X in the table indicates that option is enabled. An X enclosed in parentheses (X) indicates that the option can be on or off without impacting the action performed during the mirror.

Not all job types have the source newer option available.

Source Server Properties	Job Properties			Action Performed
Block Checksum Option	File Differences Option	Source Newer Option	Block Checksum Option	
(X)	X			Any file that is different on the source and target based on the date, time, size, and/or attribute is transmitted to the target. The mirror sends the entire file.
(X)	X	X		Any file that is newer on the source than on the target based on date and/or time is transmitted to the target. The mirror sends the entire file.
	X		X	Any file that is different on the source and target based on date, time, size, and/or attributed is flagged as different. The mirror then performs a checksum comparison on the flagged files and only sends those blocks that are different.
X	X		X	The mirror performs a checksum comparison on all files and only sends those blocks that are different.
(X)	X	X	X	Any file that is newer on the source than on the target based on date and/or time is flagged as different. The mirror then performs a checksum comparison on the flagged files and only sends those blocks that are different.

- **Enable scheduled verification**—Verification is the process of confirming that the source replica data on the target is identical to the original 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, can automatically initiate a remirror, if configured. The remirror ensures data integrity between the source and target. When this option is enabled, Double-Take will verify the source replica data on the target and generate a verification log.

  

  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.

  - **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 at the specified date and time.
  - **Mirror files to the target server when verifying**—When this option is enabled, in addition to verifying the data and generating a log, Double-Take will also mirror to the target any protected files that are different on the source.
    - **Mirror only if the file on the source is newer than the copy on the target**—This option is not available for full server to ESX jobs.
    - **Use block checksum for comparisons**—For those files flagged as different, the mirroring process can perform a block checksum comparison and send only those blocks that are different.
- **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 the calculation is enabled, it is completed before the job starts mirroring, which can take a significant amount of time depending on the number of files and system performance. If your job contains a large number of files, for example, 250,000 or more, you may want to disable the calculation so that data will start being mirrored sooner. 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 to ESX jobs. The calculation time is when the system state protection processes hard links. If you disable the calculation, the hard link processing will not occur and you may have problems after failover, especially if your source is Windows 2008.

- **Delete orphaned files**—An orphaned file is a file that exists in the replica data on the

target, but does not exist in the protected data on the source. This option specifies if orphaned files should be deleted on the target during a mirror, verification, or restoration.

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

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 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.

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 95 for more information.

During a mirror, orphaned file processing success messages will be logged to a separate orphaned file log. 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 100 for details on the location of that folder. The orphaned log file is overwritten during each orphaned file processing during a mirror, and the log file will be a maximum of 50 MB.

Network Route

For **Send data to the target server using this route**, Double-Take will select, by default, a target route for transmissions. If desired, specify an alternate route on the target that the data will be transmitted through. This allows you to select a different route for Double-Take traffic. For example, you can separate regular network traffic and Double-Take traffic on a machine with multiple IP addresses.

The IP address used on the source will be determined through the Windows route table.

Snapshots

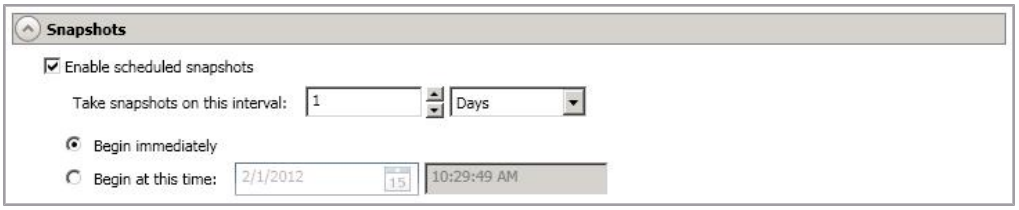

A snapshot is an image of the source replica data on the target taken at a single point in time. You can failover to a snapshot. However, you cannot access the snapshot to recover specific files or folders.

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 starting at a later date and time. Specify the date and time parameters to indicate when you want to start.

See *Managing snapshots* on page 161 for details on taking manual snapshots and deleting snapshots.

Snapshots are stored on the source replica on the target, so be sure that you configure your replica virtual machine disks large enough to maintain snapshots. Also, you may want to set the size limit on how much space snapshots can use. See your VSS documentation for more details on setting a size limit.

Compression

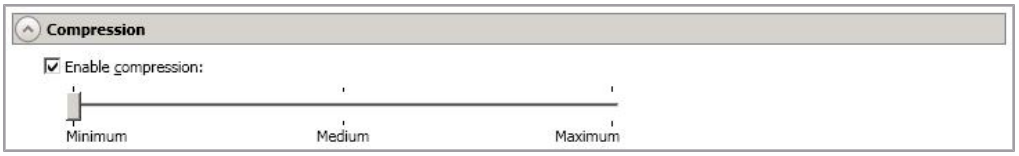

To help reduce the amount of bandwidth needed to transmit Double-Take data, compression allows you to compress data prior to transmitting it across the network. In a WAN environment this provides optimal use of your network resources. If compression is enabled, the data is compressed before it is transmitted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

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

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

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

Bandwidth

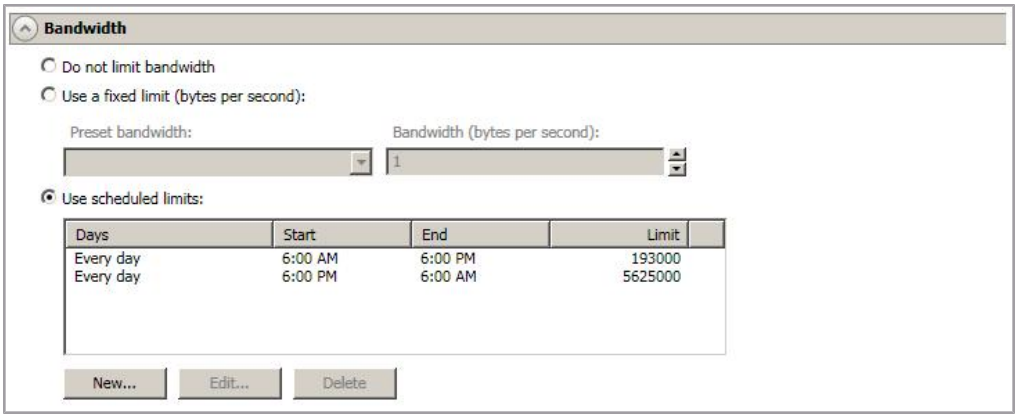

Bandwidth limitations are available to restrict the amount of network bandwidth used for Double-Take data transmissions. When a bandwidth limit is specified, Double-Take never exceeds that allotted amount. The bandwidth not in use by Double-Take is available for all other network traffic.

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

The scheduled option is not available if your source is a cluster.

- **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**.

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

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

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

Managing and controlling full server to ESX jobs

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

The jobs displayed in the right pane depend on the server group folder selected in the left pane. Every job for each server in your console session is displayed when the **Jobs on All Servers** group is selected. If you have created and populated server groups (see *Managing servers* on page 65), then only the jobs associated with the server or target servers in that server group will be displayed in the right pane.

- See *Overview job information displayed in the top pane* on page 471
- See *Detailed job information displayed in the bottom pane* on page 473
- See *Job controls* on page 475

Overview job information displayed in the top pane

The top pane displays high-level overview information about your jobs.

Column 1 (Blank)

The first blank column indicates the state of the job.

 The job is in a healthy state.

 The job is in a warning state. This icon is also displayed on any server groups that you have created that contain a job in a warning state.

 The job is in an error state. This icon is also displayed on any server groups that you have created that contain a job in an error state.

 The job is in an unknown state.

Job

The name of the job

Source Server

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

Target Server

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

Job Type

Each job type has a unique job type name. This job is a Full server to ESX job. For a complete list of all job type names, press F1 to view the Double-Take Console online help.

Activity

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

Mirror Status

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

Replication Status

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

Transmit Mode

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

Detailed job information displayed in the bottom pane

The details displayed in the bottom pane of the **Manage Jobs** page provide additional information for the job highlighted in the top pane. If you select multiple jobs, the details for the first selected job will be displayed.

Name

The name of the job

Target data state

- **OK**—The data on the target is in a good state.
- **Mirroring**—The target is in the middle of a mirror process. The data will not be in a good state until the mirror is complete.
- **Mirror Required**—The data on the target is not in a good state because a remirror is required. This may be caused by an incomplete or stopped mirror or an operation may have been dropped on the target.
- **Restore Required**—The data on the source and target do not match because of a failover condition. Restore the data from the target back to the source. If you want to discard the changes on the target, you can remirror to resynchronize the source and target.
- **Snapshot Reverted**—The data on the source and target do not match because a snapshot has been applied on the target. Restore the data from the target back to the source. If you want to discard the changes on the target, you can remirror to resynchronize the source and target.
- **Busy**—The source is low on memory causing a delay in getting the state of the data on the target.
- **Not Loaded**—Double-Take target functionality is not loaded on the target server. This may be caused by an activation code error.
- **Unknown**—The console cannot determine the status.

Mirror remaining

The total number of mirror bytes that are remaining to be sent from the source to the target

Mirror skipped

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

Replication queue

The total number of replication bytes in the source queue

Disk queue

The amount of disk space being used to 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**.

Connected since

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

Recent activity

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

Additional information

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

Job controls

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

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

Create a New Job 

This button leaves the **Manage Jobs** page and opens the **Get Started** page.

View Job Details 

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

Delete 

Stops (if running) and deletes the selected jobs.

If you no longer want to protect the source and no longer need the replica of the source on the target, select to delete the associated replica virtual machine. Selecting this option will remove the job and completely delete the replica virtual machine on the target.

If you no longer want to mirror and replicate data from the source to the target but still want to keep the replica of the source on the target, select to keep the associated replica virtual machine. You may want to use this option to relocate the virtual hard disks and create a new job between the original source and the new location. Selecting this option, will preserve the source replica on the target.

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 77. You will remain on the **Manage Jobs** page after updating the server credentials. If your servers use the same credentials, make sure you also update the credentials on the **Manage Servers** page so that the Double-Take Console can authenticate to the servers in the console session. See *Managing servers* on page 65.

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. Failover monitoring will continue 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. Failover monitoring will continue while the job is stopped.

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 161 for more information.

Failover or Cutover 

Starts the failover process. See *Failing over full server to ESX jobs* on page 489 for the process and details of failing over a full server to ESX job.

Failback 

Starts the failback process. Failback does not apply to full server to ESX jobs.

Restore 

Starts the restoration process. Restore does not apply to full server to ESX jobs.

Reverse 

Reverses protection. Reverse protection does not apply to full server to ESX jobs.

Undo Failover 

Cancels a test failover by undoing it. This resets the servers and the job back to their original state. See *Failing over full server to ESX jobs* on page 489 for the process and details of undoing a failed over full server to ESX job.

View Job Log 

Opens the job log. On the right-click menu, this option is called **View Logs**, and you have the option of opening the job log, source server log, or target server log. See *Viewing the log files through the Double-Take Console* on page 678 for details on all three of these logs.

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 all files**—All protected files will be mirrored from the source to the target.
  - **Mirror different files**—Only those protected files that are different based on date and time, size, or attributes will be mirrored from the source to the target.
    - **Mirror only if the file on the source is newer than the copy on the target**—This option is available for full server to ESX jobs, but

ideally it should not be used.

- **Use block checksum for comparisons**—For those files flagged as different, the mirroring process can perform a block checksum comparison and send only those blocks that are different.

- **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 the calculation is enabled, it is completed before the job starts mirroring, which can take a significant amount of time depending on the number of files and system performance. If your job contains a large number of files, for example, 250,000 or more, you may want to disable the calculation so that data will start being mirrored sooner. 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 to ESX jobs. The calculation time is when the system state protection processes hard links. If you disable the calculation, the hard link processing will not occur and you may have problems after failover, especially if your source is Windows 2008.

- **Verify**—Even if you have scheduled the verification process, you can run it manually any time a mirror is not in progress.
  - **Create verification report only**—This option verifies the data and generates a verification log, but it does not remirror any files that are different on the source and target. See *Verification log* on page 102 for details on the log file.
  - **Mirror files to the target server automatically**—When this option is enabled, in addition to verifying the data and generating a log, Double-Take will also mirror to the target any protected files that are different on the source.
    - **Mirror only if the file on the source is newer than the copy on the target**—This option is available for full server to ESX jobs, but ideally it should not be used.
    - **Use block checksum for comparisons**—For those files flagged as different, the mirroring process can perform a block checksum comparison and send only those blocks that are different.
- **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.

- **Update Shares**—Shares are not applicable because they are automatically included with the system state that is being protected with the entire server.

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 65.

Type a server name

Displays only jobs that contain the text you entered. If you have created and populated server groups, then only jobs that contain the text you entered associated with the server or target servers in that server group will be displayed. See *Managing servers* on page 65.

Overflow Chevron 

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

Viewing full server to ESX job details

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

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

Job name

The name of the job

Job type

Each job type has a unique job type name. This job is a Full server to ESX job. For a complete list of all job type names, press F1 to view the Double-Take Console online help.

Health

 The job is in a healthy state.

 The job is in a warning state.

 The job is in an error state.

 The job is in an unknown state.

Activity

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

Connection ID

The incremental counter used to number connections. The number is incremented when a connection is created. It is also incremented by internal actions, such as an auto-disconnect and auto-reconnect. The lowest available number (as connections are created, stopped, deleted, and so on) will always be used. The counter is reset to one each time the Double-Take service is restarted.

Transmit mode

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

Target data state

- **OK**—The data on the target is in a good state.
- **Mirroring**—The target is in the middle of a mirror process. The data will not be in a good state until the mirror is complete.
- **Mirror Required**—The data on the target is not in a good state because a remirror is required. This may be caused by an incomplete or stopped mirror or an operation may have been dropped on the target.
- **Restore Required**—The data on the source and target do not match because of a failover condition. Restore the data from the target back to the source. If you want to discard the changes on the target, you can remirror to resynchronize the source and target.
- **Snapshot Reverted**—The data on the source and target do not match because a snapshot has been applied on the target. Restore the data from the target back to the source. If you want to discard the changes on the target, you can remirror to resynchronize the source and target.
- **Busy**—The source is low on memory causing a delay in getting the state of the data on the target.
- **Not Loaded**—Double-Take target functionality is not loaded on the target server. This may be caused by an activation code error.
- **Unknown**—The console cannot determine the status.

Target route

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

Compression

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

Encryption

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

Bandwidth limit

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

Connected since

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

Additional information

Depending on the current state of your job, you may see additional information displayed to keep you informed about the progress and status of your job. If there is no

additional information, you will see (None) displayed.

Mirror status

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

Mirror percent complete

The percentage of the mirror that has been completed

Mirror remaining

The total number of mirror bytes that are remaining to be sent from the source to the target

Mirror skipped

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

Replication status

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

Replication queue

The total number of replication bytes in the source queue

Disk queue

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

Bytes sent

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

Bytes sent compressed

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

Validating a full server to ESX job

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

1. From the **Manage Jobs** page, highlight the job and click **View Job Details** in the toolbar.
2. In the **Tasks** area on the right on the **View Job Details** page, click **Validate job properties**.
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 enable protection. Depending on the error, you may be able to click **Fix** or **Fix All** and let Double-Take correct the problem for you. For those errors that Double-Take cannot correct automatically, you will need to modify the source or target to correct the error, or you can select a different target. You must revalidate the selected servers, by clicking **Recheck**, until the validation check passes without errors.

   Validation checks for an existing job are logged to the job log on the target server. See *Log files* on page 677 for details on the various log files.
4. Once your servers have passed validation, click **Close**.

Editing a full server to ESX job

Use these instructions to edit a full server to ESX job.

1. From the **Manage Jobs** page, highlight the job and click **View Job Details** in the toolbar.
2. In the **Tasks** area on the right on the **View Job Details** page, click **Edit job properties**. (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 for your full server to ESX job as when you created the job, but you will not be able to edit all of them. If desired, edit those options that are configurable for an existing job. See *Creating a full server to ESX job* on page 445 for details on each job option.

   

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

4. If you want to modify the workload items or replication rules for the job, click **Edit workload or replication rules**. Modify the **Workload item** you are protecting, if desired. Additionally, you can modify the specific **Replication Rules** for your job.

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

   You can also enter wildcard rules, however you should do so carefully. Rules are applied to files that are closest in the directory tree to them. If you have rules that include multiple folders, an exclusion rule with a wild card will need to be added for each folder that it needs applied to. For example, if you want to exclude all .log files from D:\ and your rules include D:\, D:\Dir1 , and D:\Dir2, you would need to add the exclusion rule for the root and each subfolder rule. So you will need to add exclude rules for D:*.log , D:\Dir1*.log, and D:\Dir2*.log.

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

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

   Errors are designated by a white X inside a red circle. Warnings are designated by a black exclamation point (!) inside a yellow triangle. A successful validation is designated by a white checkmark inside a green circle. You can sort the list by the icon to see errors, warnings, or successful validations together. Click on any of the validation items to see details. You must

correct any errors before you can enable protection. Depending on the error, you may be able to click **Fix** or **Fix All** and let Double-Take correct the problem for you. For those errors that Double-Take cannot correct automatically, you will need to modify the source or target to correct the error, or you can select a different target. You must revalidate the selected servers, by clicking **Recheck**, until the validation check passes without errors.

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.

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

Viewing a full server to ESX job log

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

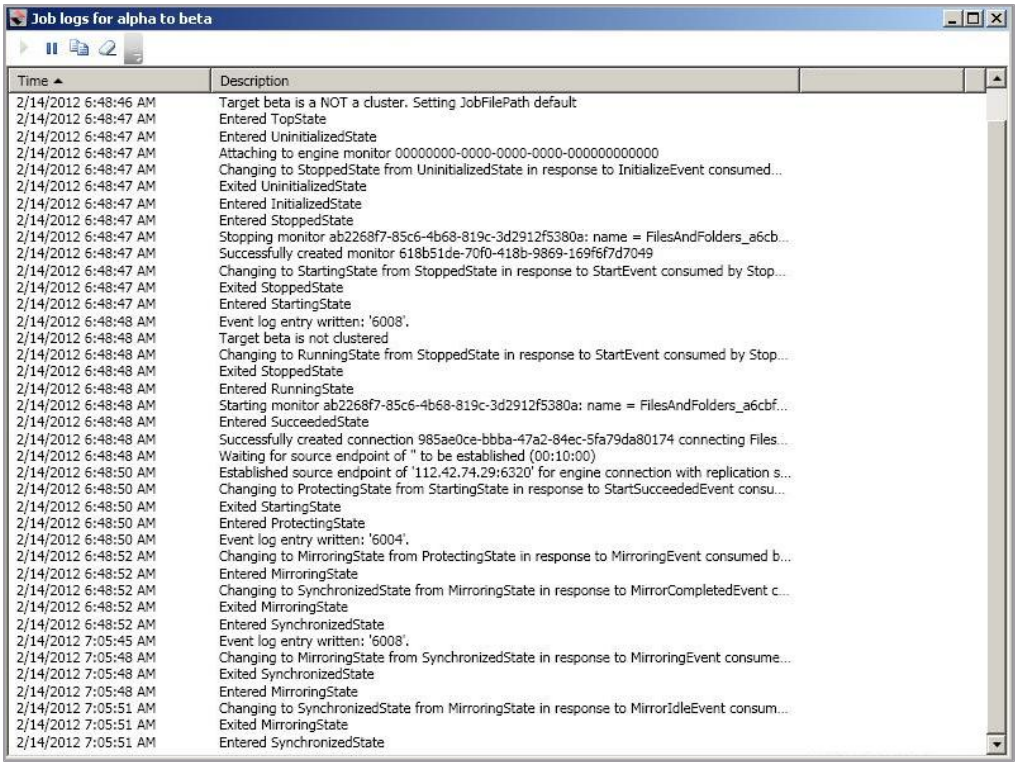

The following table identifies the controls and the table columns in the **Job logs** window.

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.

Failing over full server to ESX jobs

When a failover condition has been met, failover will be triggered automatically if you disabled the wait for user option during your failover configuration. If the wait for user before failover option is enabled, you will be notified in the console when a failover condition has been met. At that time, you will need to trigger it manually from the console when you are ready.

1. On the **Manage Jobs** page, highlight the job that you want to failover and click **Failover or Cutover** in the toolbar.
2. Select the type of failover to perform.
    - **Failover to live data**—Select this option to initiate a full, live failover using the current data on the target. This option will shutdown the source machine (if it is online), stop the protection job, and start the replica virtual machine on the target with full network connectivity.
    - **Perform test failover**—Select this option to perform a test failover using the current data on the target. This option will leave the source machine online, stop the protection job, and start the replica virtual machine on the target without network connectivity.
    - **Failover to a snapshot**—Select this option to initiate a full, live failover without using the current data on the target. Instead, select a snapshot and the data on the target will be reverted to that snapshot. This option will not be available if there are no snapshots on the target or if the target does not support snapshots. This option is also not applicable to clustered environments. To help you understand what snapshots are available, the **Type** indicates the kind of snapshot.
        - **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. Select how you want to handle the data in the target queue. You may want to check the amount of data in queue on the target by reviewing the *Statistics* on page 688 or *Performance Monitor* on page 790.
    - **Apply data in target queues before failover or cutover**—All of the data in the target queue will be applied before failover begins. The advantage to this option is that all of the data that the target has received will be applied before failover begins. 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 failover or cutover immediately**—All of the data in the target queue will be discarded and failover will begin immediately. The advantage to this option is that failover will occur immediately. The disadvantage is that any data in the target queue will be lost.
    - **Revert to last good snapshot if target data state is bad**—If the target data is in a bad state, Double-Take will automatically revert to the last good Double-Take snapshot before failover begins. If the target data is in a good state, Double-Take will not revert the target data. Instead, Double-Take will apply the data in the target queue and then failover. The advantage to this option is that good data on the target is guaranteed to be used. The

disadvantage is that if the target data state is bad, you will lose any data between the last good snapshot and the failure.

4. When you are ready to begin failover, click **Failover**.

Because the Windows product activation is dependent on hardware, you may need to reactivate your Windows registration after failover. In most cases when you are using Windows 2003, you can follow the on-screen prompts to complete the reactivation. However, when you are using Windows 2008, the reactivation depends on several factors including service pack level, Windows edition, and your licensing type. If a Windows 2008 target comes online after failover 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.

Depending on your replica configuration, you may have to reboot your replica after failover. You will be prompted to reboot if it is necessary.

After failover, if you attempt to use a full server job to revert back to your original configuration, you will need to perform a few additional tasks before creating the full server job. Contact technical support if you need assistance with these steps.

1. On either the source or target, stop the Double-Take and Double-Take Management Service services.

2. Remove the GUID value from HKEY_LOCAL_MACHINE\SOFTWARE\NSI Software\Double-Take\CurrentVersion\Communication\UniqueId. Do not delete the UniqueId key. Only delete the GUI value within the key.
3. Restart the the Double-Take and Double-Take Management Service services.
4. Remove and then add your servers back into the Double-Take Console.
5. Install a different license on the original source and complete a host transfer if necessary.

If your source was using vCenter, you may have problems with failover if vCenter is down or if it is unreachable. See the technical support article 34768 for details on how to complete failover in this situation.

---

5. If you performed a test failover, you can undo it by selecting **Undo Failover or Cutover** in the toolbar. The replica virtual machine on the target will be shut down and the protection job will be restarted performing a file differences mirror.
6. There is no reverse or failback once you have failed over. If you need to go back to your original hardware, delete the job and re-create a new one if your original source was a virtual server. If your original source was a physical server, you will need to use a full server job.

Chapter 13 V to ESX protection

Create a V to ESX job when you want to protect a virtual machine on an ESX host to another ESX host. This protection type is at the virtual machine guest level. You can reverse the protection after failover has occurred.

- See *V to ESX requirements* on page 493—V to ESX protection includes specific requirements for this type of protection.
- See *Creating a V to ESX job* on page 496—The section includes step-by-step instructions for creating a V to ESX job.
- See *Managing and controlling V to ESX jobs* on page 504—You can view status information about your V to ESX jobs and learn how to control these jobs.
- See *Failing over V to ESX jobs* on page 541—Use this section when a failover condition has been met or if you want to failover manually.
- See *Reversing V to ESX jobs* on page 543—Use this section to reverse protection. The source replica on the target is now sending data back to the original source.

V to ESX requirements

In addition to the *Core Double-Take requirements* on page 23, use these requirements for V to ESX protection.

- **Source virtual server**—The source virtual server can be any of the operating systems listed in the *Core Double-Take requirements* on page 23. The source server does not need Double-Take installed on it. If it is not, the installation will be completed during the protection process, however the virtual must meet the Windows firewall requirements as described in the *Core Double-Take requirements* on page 23.

  

  In order for the installation to be completed during the protection process, the Double-Take Console should have enough valid licenses in the license inventory for each of the virtual machines you are protecting. If you do not have enough valid licenses available, you will have to install Double-Take on each virtual machine individually and then restart each protection job. See *Managing the Double-Take license inventory* on page 49.

- **Source and target ESX server**—The ESX server that will host your source and target must be the same version of ESX. The version of ESX can be any of the following.
  - ESX 4.0.x or 4.1 Standard, Advanced, Enterprise, or Enterprise Plus
  - ESXi 4.0.x or 4.1 Standard, Advanced, Enterprise, or Enterprise Plus
  - ESXi 5.0 Standard, Enterprise, or Enterprise Plus
  - ESXi 5.1 Essentials, Essentials Plus, Standard, Enterprise, or Enterprise Plus
  - ESXi 5.5 Essentials, Essentials Plus, Standard, Enterprise, or Enterprise Plus

  ---

  

  If you are using the Standard edition of ESX 4.0 or ESXi 4.0, you must have update 1 or later.

  If your source is Windows 2008 R2, your ESX server must have ESX 4.0 update 1 or later.

  If your source is Windows 2012 or 2012 R2, your ESX server must have ESXi 5.0 update 1 or later.

- **vCenter**—vCenter is not required, but if you are using it, then you must use version 4.1 or later.
- **vMotion**—Host vMotion is only supported if you are using vCenter. Storage vMotion is not supported.
- **Virtual recovery appliance**—The target ESX server must have an existing virtual machine, known as a virtual recovery appliance, that meets the following requirements. (When you establish protection, the virtual recovery appliance will create new virtual servers, mount disks, format disks, and so on. If failover occurs, the new virtual machines are detached from the virtual recovery appliance and powered on. Once the new virtual machines are online, they will have the identity, data, and system state of their source. Since the virtual recovery appliance maintains its

own identity, it can be reused for additional failovers.) You can optionally have an existing virtual recovery appliance on the source ESX server to reverse your protection if desired.

- **Operating system**—The virtual recovery appliance can be any of the operating systems listed in the *Core Double-Take requirements* on page 23.
- **Operating system version**—The virtual recovery appliance must have the same or newer operating system than the source (not including service pack level).
- **Operating system installation location**—Because VMware boots from the first bootable volume that is discovered, the operating system must be installed to SCSI device 0, Slot 0 on the virtual recovery appliance.
- **Double-Take**—The virtual recovery appliance must have Double-Take installed and licensed on it.
- **Microsoft .NET Framework**—Microsoft .NET Framework version 3.5 Service Pack 1 is required on the virtual recovery appliance. This version is not included in the .NET version 4.0 release. Therefore, even if you have .NET version 4.0 installed, you will also need version 3.5.1. If you are using Windows 2008 or earlier, you can install this version from the Double-Take DVD, via a web connection during the Double-Take installation, or from a copy you have obtained manually from the Microsoft web site.If you are using Windows 2008 R2 or later, you can enable it through Windows features.
- **Snapshots**—Do not take snapshots of the virtual recovery appliance, because they will interfere with proper failover.

- **Disk controller**—VMware Paravirtual SCSI Controllers are not supported.
- **IP addressing**—IPv4 is the only supported IP version.
- **Domain controllers**—If you are protecting a domain controller, it will start in a non-authoritative restore mode after failover. This means that if it 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 this is the only domain controller, this is not an issue.
- **Snapshots**—Double-Take snapshots are not supported with V to ESX jobs.
- **Supported configurations**—The following table identifies the supported configurations for a V to ESX job. (These are source and target configurations for the host, not the virtual machines.)

Configuration		Supported	Not Supported
Source to target configuration[1]	One to one, active/standby	X	
	One to one, active/active	X	
	Many to one	X	
	One to many	X	
	Chained		X
	Single server	X	

Configuration		Supported	Not Supported
Server configuration[2]	Standalone to standalone	X	
	Standalone to cluster	X	
	Cluster to standalone	X	
	Cluster to cluster	X	
Upgrade configuration[3]	Upgrade 5.3 Windows ESX to ESX job to 7.0 V to ESX job	X	
	Upgrade 6.0 V to ESX job to 7.0 V to ESX job	X	
Version 7.0 console[4,5]	Version 7.0 console can create job for 5.3 source and 5.3 target		X
	Version 7.0 console can create job for 6.0 source and 6.0 target	X	
	Version 7.0 console can create job for 7.0 source and 7.0 target	X	

1. See *Supported configurations* on page 16 for details on each of the source to target configurations.
2. Supported cluster configurations can be used to a virtual machine, however, there is no cluster resource or cluster-awareness with this configuration.
3. When upgrading from version 5.3, you can perform a rolling upgrade where you update the target server first. After the upgrade is complete, the source will automatically reconnect to the target. At this point, the job will be an unmanaged job that you can delete or failover. No other job controls will be available. Once you upgrade you source, the job will be fully controllable.
4. Once you upgrade your console to version 7.0, existing jobs that are running version 5.3 will not appear in the console until the target of the job is upgraded to version 7.0.
5. Newer job options available in the version 7.0 console will not be functional when creating jobs for servers running version 6.0.

Creating a V to ESX job

Use these instructions to create a V to ESX job.

1. Click **Get Started** from the toolbar.
2. Select **Double-Take Availability** and click **Next**.
3. Select **Protect multiple Windows virtual servers on ESX with a Windows Virtual Recovery Appliance** and click **Next**.
4. Choose your source server. This is the ESX host that contains the virtual servers that you want to protect.

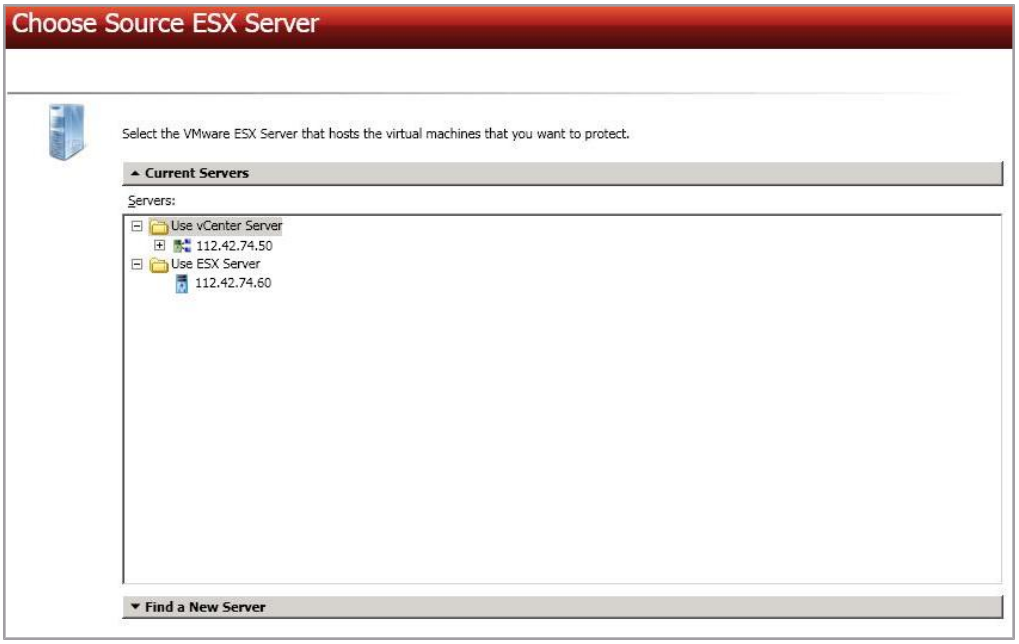

   - **Current Servers**—This list contains the vCenter and ESX servers currently available in your console session. Select your ESX server from the list.
   - **Find a New Server**—If the server you need is not in the **Current Servers** list, click the **Find a New Server** heading.
     - **vCenter Server**—Select your vCenter server from the list. If your vCenter server is not in the list, click **Add VirtualCenter Server**, specify the server and valid credentials, and click **Add**. If you are not using vCenter, select **None**.
     - **ESX Server**—Specify the name or IP address of the ESX server.
     - **User name**—This field will only be available if you are not using vCenter. In this case, specify the root user or another user that has the administrator role on the specified ESX server.
     - **Password**—Specify the password associated with the **User name** you entered.
     - **Domain**—If you are working in a domain environment, specify the Domain.
5. Click **Next** to continue.
6. Select the virtual machines on your source that you want to protect. Virtual machines that are

suspended or powered off, or those that do not have a valid operating system according to ESX will not be available for selection. A separate job will be created for each source that you select for protection.

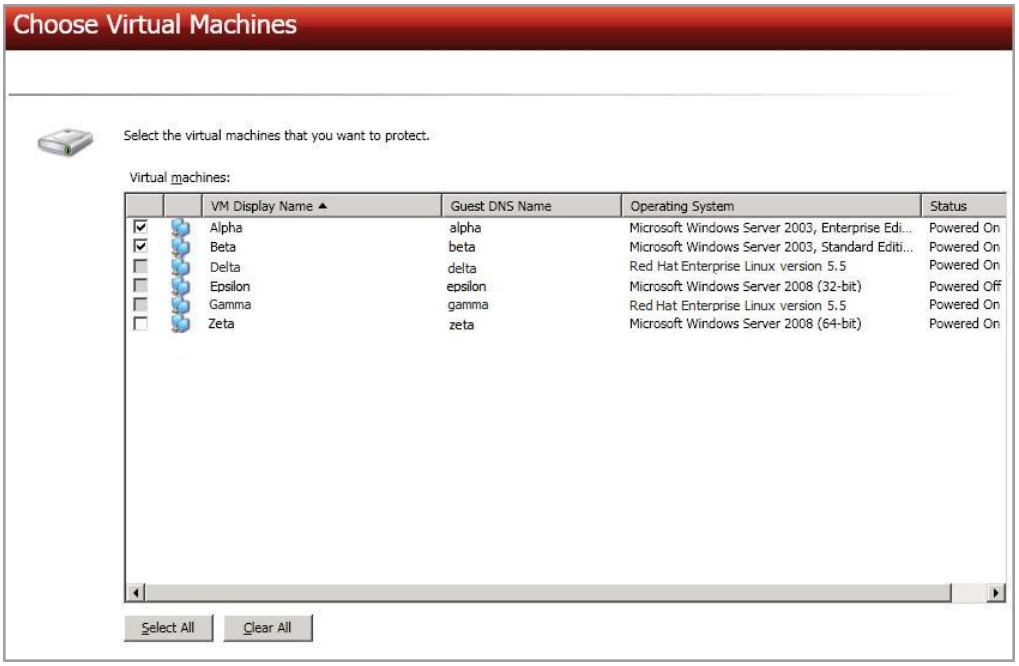

Each source can have no more than ten NICs enabled.

7. Click **Next** to continue.
8. Choose the ESX server where your target virtual recovery appliance is located. This is also the server where your replica virtual machine will be located.

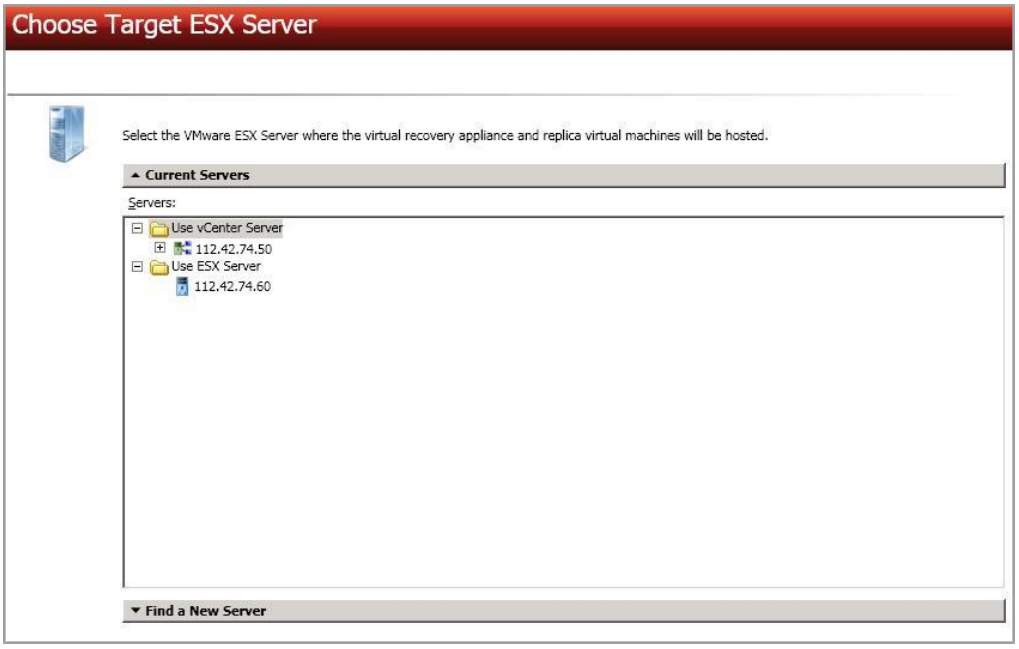

- **Current Servers**—This list contains the vCenter and ESX servers currently available in your console session. Select your ESX server from the list.
- **Find a New Server**—If the server you need is not in the **Current Servers** list, click the **Find a New Server** heading.
  - **vCenter Server**—Select your vCenter server from the list. If your vCenter server is not in the list, click **Add VirtualCenter Server**, specify the server and valid credentials, and click **Add**. If you are not using vCenter, select **None**.
  - **ESX Server**—Specify the name or IP address of the ESX server.
  - **User name**—This field will only be available if you are not using vCenter. In this case, specify the root user or another user that has the administrator role on the specified ESX server.
  - **Password**—Specify the password associated with the **User name** you entered.
  - **Domain**—If you are working in a domain environment, specify the Domain.

9. Click **Next** to continue.
10. Choose your virtual recovery appliance on your target ESX server. Only valid virtual recovery appliances will be displayed in the list. If necessary, click **Provide Credentials** and specify a valid user on the virtual recovery appliance you have selected. The user must be a member of the local Double-Take Admin and local administrators security groups. See *V to ESX requirements* on page 493 for details on the virtual recovery appliance.

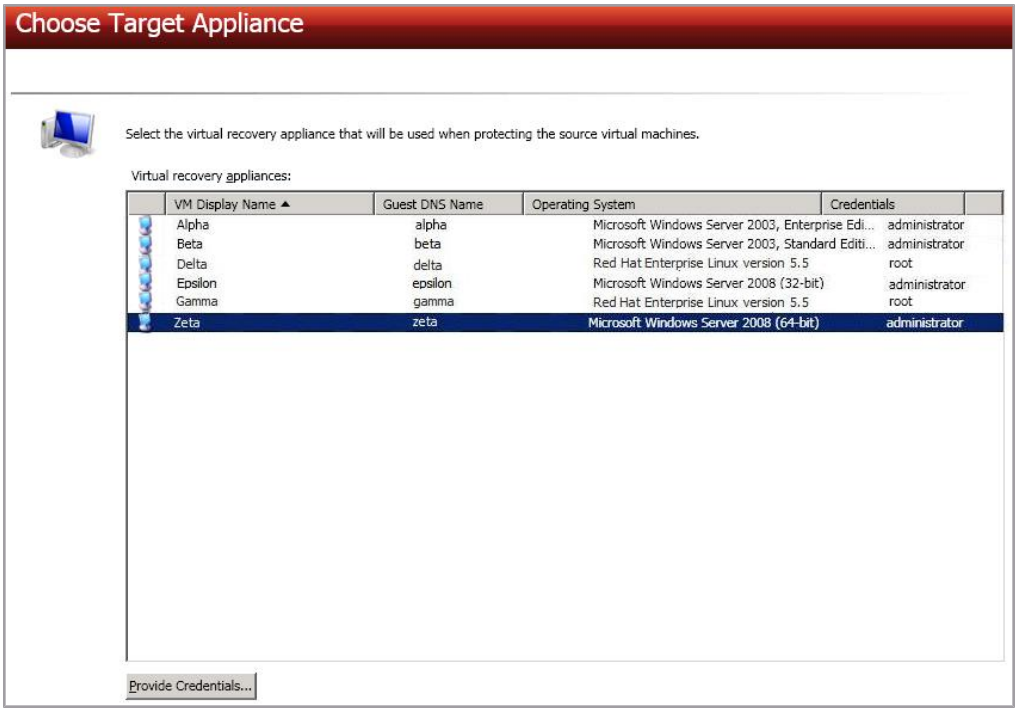

11. Click **Next** to continue.
12. If desired, select a virtual recovery appliance on your source ESX server. This appliance will be used during the reverse process, allowing you to protect the source replica virtual server on the target back to the original source ESX host. If necessary, click **Provide Credentials** and specify a valid user on the virtual recovery appliance you have selected. The user must be a member of the Double-Take Admin security group. You can skip this step and select a reverse appliance after the job has been created, however, you must have a reverse appliance selected before failover. See *V to ESX requirements* on page 493 for details on the virtual recovery appliance.

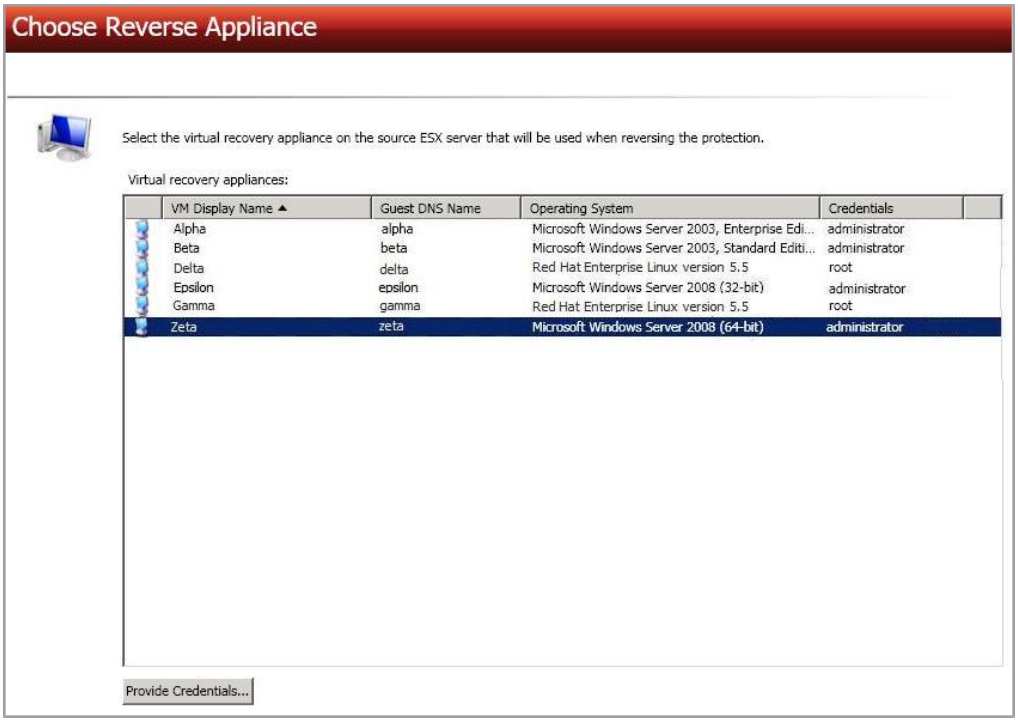

13. Click **Next** to continue.
14. Select your protection options.

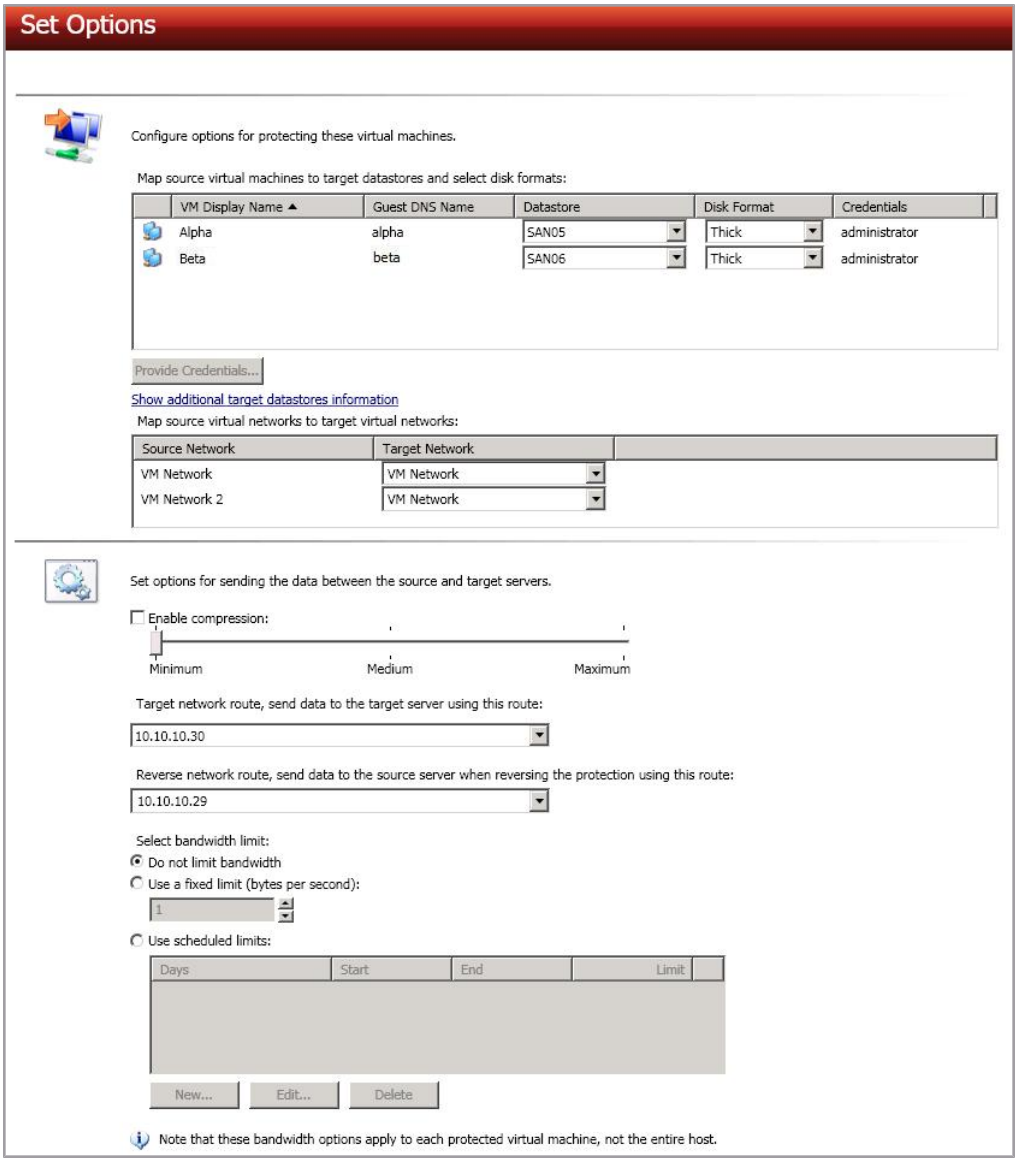

- **Datastore**—For each virtual server you are protecting, select a datastore on the target ESX host where the replica source virtual server will be created. If necessary, click **Provide Credentials** and specify a valid user on each virtual server. If you want to view disk information for the datastores on your target ESX host, click **Show additional target datastores information**. To hide the information, click **Hide additional target datastores information**.
- **Disk Format**—For each virtual server you are protecting, specify the type of disk, **Flat** (for ESX 5 only), **Thick** or **Thin**, that will be created on the replica virtual server.
- **Target Network**—For each virtual network on the source ESX host, select a network on the target ESX host to handle networking for the virtual servers after failover.
- **Enable 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. On a default Double-Take installation, compression is disabled. To enable it, select this option and 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.

- **Target network route**—By default, Double-Take will select a default target route for transmissions. If desired, select a different target route. 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.
- **Reverse network route**—If you selected a reverse appliance, Double-Take will select a default source route for transmissions. If desired, select a different source route.
- **Select bandwidth limit**—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**—Enter a value, in bytes per second, to limit data transmission. This is the maximum amount of data that will be transmitted per second.

- **Use scheduled limits**—Use a schedule to limit bandwidth for different times. Schedules that you create will be maintained if you change to a fixed limit or to no limit.
  - **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 28000 bytes per second.
  - **Edit**—Click **Edit** to modify an existing scheduled bandwidth limit.
  - **Delete**—Click **Delete** to remove a scheduled bandwidth limit.

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

    Errors are designated by a white X inside a red circle. Warnings are designated by a black exclamation point (!) inside a yellow triangle. A successful validation is designated by a white checkmark inside a green circle. You can sort the list by the icon to see errors, warnings, or successful validations together. Click on any of the validation items to see details. You must correct any errors before you can enable protection.

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

Because this job type allowed you to create more than one job at a time, options that are job specific were not presented during the job creation process. Once the jobs are created, you can edit each one individually to see and edit all of the various job options. See *Editing a V to ESX job* on page 518.

Managing and controlling V to ESX jobs

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

The jobs displayed in the right pane depend on the server group folder selected in the left pane. Every job for each server in your console session is displayed when the **Jobs on All Servers** group is selected. If you have created and populated server groups (see *Managing servers* on page 65), then only the jobs associated with the server or target servers in that server group will be displayed in the right pane.

- See *Overview job information displayed in the top pane* on page 504
- See *Detailed job information displayed in the bottom pane* on page 506
- See *Job controls* on page 508

Overview job information displayed in the top pane

The top pane displays high-level overview information about your jobs.

Column 1 (Blank)

The first blank column indicates the state of the job.

 The job is in a healthy state.

 The job is in a warning state. This icon is also displayed on any server groups that you have created that contain a job in a warning state.

 The job is in an error state. This icon is also displayed on any server groups that you have created that contain a job in an error state.

 The job is in an unknown state.

Job

The name of the job

Source Server

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

Target Server

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

Job Type

Each job type has a unique job type name. This job is a V to ESX job. For a complete list of all job type names, press F1 to view the Double-Take Console online help.

Activity

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

Mirror Status

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

Replication Status

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

Transmit Mode

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

Detailed job information displayed in the bottom pane

The details displayed in the bottom pane of the **Manage Jobs** page provide additional information for the job highlighted in the top pane. If you select multiple jobs, the details for the first selected job will be displayed.

Name

The name of the job

Target data state

- **OK**—The data on the target is in a good state.
- **Mirroring**—The target is in the middle of a mirror process. The data will not be in a good state until the mirror is complete.
- **Mirror Required**—The data on the target is not in a good state because a remirror is required. This may be caused by an incomplete or stopped mirror or an operation may have been dropped on the target.
- **Restore Required**—The data on the source and target do not match because of a failover condition. Restore the data from the target back to the source. If you want to discard the changes on the target, you can remirror to resynchronize the source and target.
- **Snapshot Reverted**—The data on the source and target do not match because a snapshot has been applied on the target. Restore the data from the target back to the source. If you want to discard the changes on the target, you can remirror to resynchronize the source and target.
- **Busy**—The source is low on memory causing a delay in getting the state of the data on the target.
- **Not Loaded**—Double-Take target functionality is not loaded on the target server. This may be caused by an activation code error.
- **Unknown**—The console cannot determine the status.

Mirror remaining

The total number of mirror bytes that are remaining to be sent from the source to the target

Mirror skipped

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

Replication queue

The total number of replication bytes in the source queue

Disk queue

The amount of disk space being used to 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**.

Connected since

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

Recent activity

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

Additional information

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

Job controls

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

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

Create a New Job 

This button leaves the **Manage Jobs** page and opens the **Get Started** page.

View Job Details 

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

Delete 

Stops (if running) and deletes the selected jobs.

If you no longer want to protect the source and no longer need the replica of the source on the target, select to delete the associated replica virtual machine. Selecting this option will remove the job and completely delete the replica virtual machine on the target.

If you no longer want to mirror and replicate data from the source to the target but still want to keep the replica of the source on the target, select to keep the associated replica virtual machine. You may want to use this option to relocate the virtual hard disks and create a new job between the original source and the new location. Selecting this option, will preserve the source replica on the target.

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 77. You will remain on the **Manage Jobs** page after updating the server credentials. If your servers use the same credentials, make sure you also update the credentials on the **Manage Servers** page so that the Double-Take Console can authenticate to the servers in the console session. See *Managing servers* on page 65.

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. Failover monitoring will continue 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. Failover monitoring will continue while the job is stopped.

Take Snapshot 

Snapshots are not applicable to V to ESX jobs.

Manage Snapshots 

Snapshots are not applicable to V to ESX jobs.

Failover or Cutover 

Starts the failover process. See *Failing over V to ESX jobs* on page 541 for the process and details of failing over a V to ESX job.

Failback 

Starts the failback process. Failback does not apply to V to ESX jobs.

Restore 

Starts the restoration process. Restore does not apply to V to ESX jobs.

Reverse 

Reverses protection. The job will start mirroring in the reverse direction with the job name and log file names changing accordingly. After the mirror is complete, the job will continue running in the opposite direction. See *Reversing V to ESX jobs* on page 543 for the process and details of reversing a V to ESX job.

Undo Failover 

Cancels a test failover by undoing it. This resets the servers and the job back to their original state.. See *Failing over V to ESX jobs* on page 541 for the process and details of undoing a failed over V to ESX job.

View Job Log 

Opens the job log. On the right-click menu, this option is called **View Logs**, and you have the option of opening the job log, source server log, or target server log. See *Viewing the log files through the Double-Take Console* on page 678 for details on all three of these logs.

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 all files**—All protected files will be mirrored from the source to the target.
  - **Mirror different files**—Only those protected files that are different based on date and time, size, or attributes will be mirrored from the source to the target.
    - **Mirror only if the file on the source is newer than the copy on the target**—This option is available for V to ESX jobs, but ideally it should not be used.

- **Use block checksum for comparisons**—For those files flagged as different, the mirroring process can perform a block checksum comparison and send only those blocks that are different.

- **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 the calculation is enabled, it is completed before the job starts mirroring, which can take a significant amount of time depending on the number of files and system performance. If your job contains a large number of files, for example, 250,000 or more, you may want to disable the calculation so that data will start being mirrored sooner. 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 V to ESX jobs. The calculation time is when the system state protection processes hard links. If you disable the calculation, the hard link processing will not occur and you may have problems after failover, especially if your source is Windows 2008.

- **Verify**—Even if you have scheduled the verification process, you can run it manually any time a mirror is not in progress.
  - **Create verification report only**—This option verifies the data and generates a verification log, but it does not remirror any files that are different on the source and target. See *Verification log* on page 102 for details on the log file.
  - **Mirror files to the target server automatically**—When this option is enabled, in addition to verifying the data and generating a log, Double-Take will also mirror to the target any protected files that are different on the source.
    - **Mirror only if the file on the source is newer than the copy on the target**—This option is available for V to ESX jobs, but ideally it should not be used.
    - **Use block checksum for comparisons**—For those files flagged as different, the mirroring process can perform a block checksum comparison and send only those blocks that are different.
- **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.
- **Update Shares**—Shares are not applicable because they are automatically included with the system state that is being protected with the entire server.

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 65.

Type a server name

Displays only jobs that contain the text you entered. If you have created and populated server groups, then only jobs that contain the text you entered associated with the server or target servers in that server group will be displayed. See *Managing servers* on page 65.

Overflow Chevron 

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

Viewing V to ESX job details

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

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

Job name

The name of the job

Job type

Each job type has a unique job type name. This job is a V to ESX job. For a complete list of all job type names, press F1 to view the Double-Take Console online help.

Health

 The job is in a healthy state.

 The job is in a warning state.

 The job is in an error state.

 The job is in an unknown state.

Activity

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

Connection ID

The incremental counter used to number connections. The number is incremented when a connection is created. It is also incremented by internal actions, such as an auto-disconnect and auto-reconnect. The lowest available number (as connections are created, stopped, deleted, and so on) will always be used. The counter is reset to one each time the Double-Take service is restarted.

Transmit mode

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

Target data state

- **OK**—The data on the target is in a good state.
- **Mirroring**—The target is in the middle of a mirror process. The data will not be in a good state until the mirror is complete.
- **Mirror Required**—The data on the target is not in a good state because a remirror is required. This may be caused by an incomplete or stopped mirror or an operation may have been dropped on the target.
- **Restore Required**—The data on the source and target do not match because of a failover condition. Restore the data from the target back to the source. If you want to discard the changes on the target, you can remirror to resynchronize the source and target.
- **Snapshot Reverted**—The data on the source and target do not match because a snapshot has been applied on the target. Restore the data from the target back to the source. If you want to discard the changes on the target, you can remirror to resynchronize the source and target.
- **Busy**—The source is low on memory causing a delay in getting the state of the data on the target.
- **Not Loaded**—Double-Take target functionality is not loaded on the target server. This may be caused by an activation code error.
- **Unknown**—The console cannot determine the status.

Target route

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

Compression

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

Encryption

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

Bandwidth limit

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

Connected since

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

Additional information

Depending on the current state of your job, you may see additional information displayed to keep you informed about the progress and status of your job. If there is no

additional information, you will see (None) displayed.

Mirror status

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

Mirror percent complete

The percentage of the mirror that has been completed

Mirror remaining

The total number of mirror bytes that are remaining to be sent from the source to the target

Mirror skipped

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

Replication status

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

Replication queue

The total number of replication bytes in the source queue

Disk queue

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

Bytes sent

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

Bytes sent compressed

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

Validating a V to ESX job

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

1. From the **Manage Jobs** page, highlight the job and click **View Job Details** in the toolbar.
2. In the **Tasks** area on the right on the **View Job Details** page, click **Validate job properties**.
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 enable protection.

   Validation checks for an existing job are logged to the job log on the target server. See *Log files* on page 677 for details on the various log files.
4. Once your servers have passed validation, click **Close**.

Editing a V to ESX job

Use these instructions to edit a V to ESX job.

1. From the **Manage Jobs** page, highlight the job and click **View Job Details** in the toolbar.
2. In the **Tasks** area on the right on the **View Job Details** page, click **Edit job properties**. (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. Because you configured multiple jobs at once when you first established your protection, not all of the individual job options were available during job creation. Therefore, you will have additional job options when you edit an existing job.

   

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

   There will be additional sections on the **Edit Job Properties** page that you will be able to view only. You cannot edit those sections.

   Go to each page identified below to see the options that you can edit on the **Edit Job Properties** page. After you have configured your options, continue with the next step on page 538.

   - *General* on page 519
   - *Replica Virtual Machine Location* on page 520
   - *Replica Virtual Machine Configuration* on page 521
   - *Replica Virtual Machine Volumes* on page 522
   - *Replica Virtual Machine Network Settings* on page 523
   - *Reverse Appliance* on page 524
   - *Failover Monitor* on page 525
   - *Failover Options* on page 527
   - *Failover Identity* on page 528
   - *Mirror, Verify & Orphaned Files* on page 530
   - *Network Route* on page 534
   - *Compression* on page 535
   - *Bandwidth* on page 536

General

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

Replica Virtual Machine Location

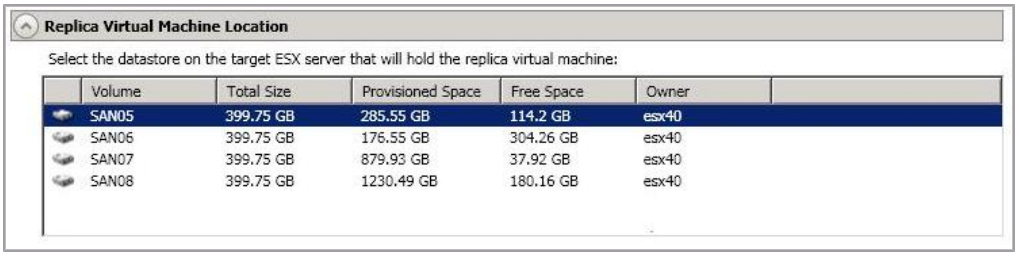

Select one of the volumes from the list to indicate the volume on the target where you want to store the configuration files for the new virtual server when it is created. The target volume must have enough **Free Space**. You can select the location of the .vmdk files under **Replica Virtual Machine Volumes**.

Replica Virtual Machine Configuration

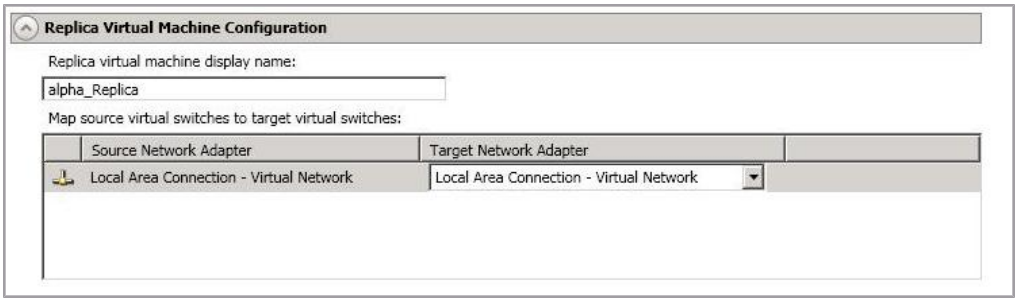

- **Replica virtual machine display name**—Specify the name of the replica virtual machine. This will be the display name of the virtual machine on the host system.
- **Map source virtual switches to target virtual switches**—Identify how you want to handle the network mapping after failover. The **Source Network Adapter** column lists the NICs from the source. Map each one to a **Target Network Adapter**, which is a virtual network on the target. You can also choose to discard the source's NIC and IP addresses.

Replica Virtual Machine Volumes

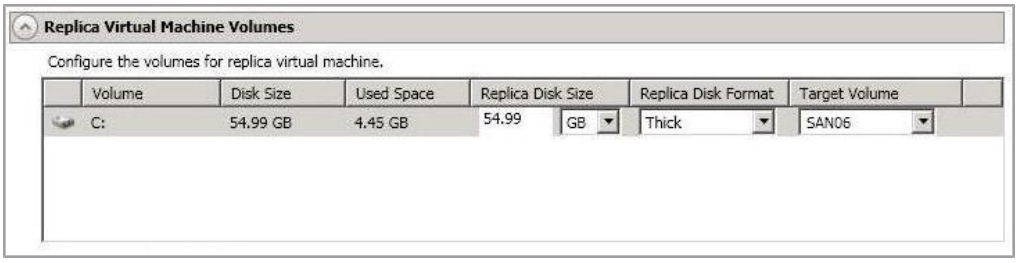

- **Replica Disk Size**—For each volume you are protecting, specify the size of the replica disk on the target. Be sure and include the value in MB or GB for the disk. The value must be at least the size of the specified **Used Space** on that volume.

  In some cases, the replica virtual machine may use more virtual disk space than the size of the source volume due to differences in how the virtual disk's block size is formatted and how hard links are handled. To avoid this issue, specify the size of your replica to be at least 5 GB larger.

  Snapshots are stored on the replica, so if you enable snapshots, be sure that you configure your replica virtual machine disk size large enough to maintain snapshots.

- **Replica Disk Format**—For each volume you are protecting, specify the format of the disk that will be created.
  - **Flat**—This disk format allocates the full amount of the disk space immediately, but does not initialize the disk space to zero until it is needed. This disk format is only available on ESX 5.
  - **Thick**—This disk format allocates the full amount of the disk space immediately, initializing all of the allocated disk space to zero.
  - **Thin**—This disk format does not allocate the disk space until it is needed.
- **Target Volume**—For each volume you are protecting, specify the volume on the target where you want to store the .vmdk files for the new replica virtual machine. You can specify the location of the virtual machine configuration files under **Replica Virtual Machine Location**.

Replica Virtual Machine Network Settings

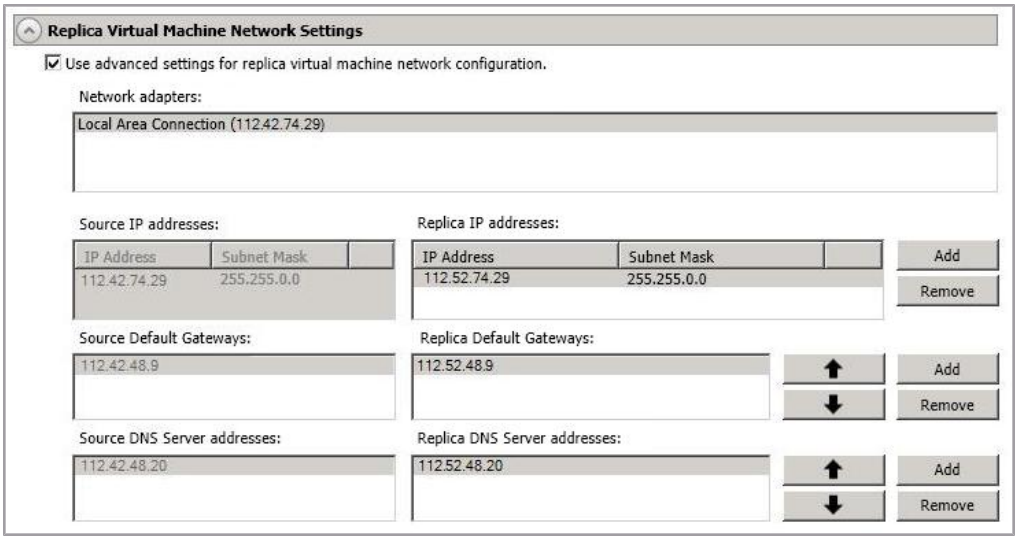

- **Use advanced settings for replica virtual machine network configuration**—Select this option to enable the replica virtual machine network setting configuration. This setting is primarily used for WAN support.
- **Network adapters**—Select a network adapter from the source and specify the **Replica IP addresses**, **Replica Default Gateways**, and **Replica DNS Server addresses** to be used after failover. If you add multiple gateways or DNS servers, you can sort them by using the arrow up and arrow down buttons. Repeat this step for each network adapter on the source.

Updates made during failover will be based on the network adapter name when protection is established. If you change that name, you will need to delete the job and re-create it so the new name will be used during failover.

If you update one of the advanced settings (IP address, gateway, or DNS server), then you must update all of them. Otherwise, the remaining items will be left blank. If you do not specify any of the advanced settings, the replica virtual machine will be assigned the same network configuration as the source.

By default, the source IP address will be included in the target IP address list as the default address. If you do not want the source IP address to be the default address on the target after failover, remove that address from the **Replica IP addresses** list.

Reverse Appliance

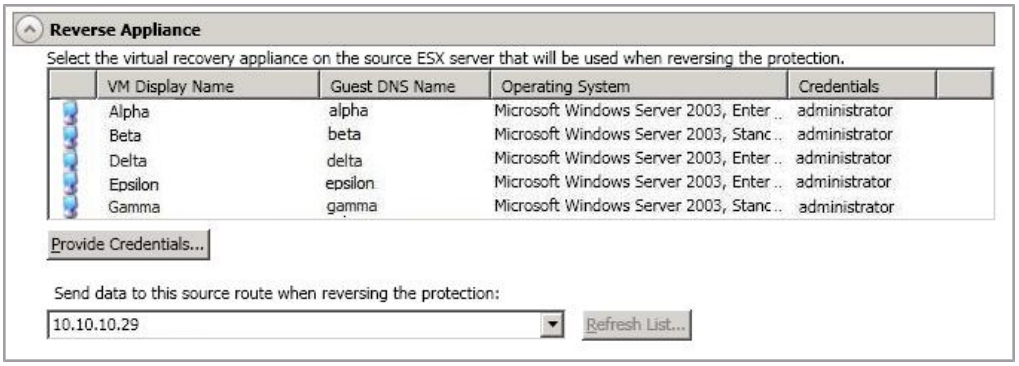

- **Select the virtual recovery appliance on the source ESX server that will be used when reversing the protection**—If desired, select a virtual recovery appliance on your source ESX server. This appliance will be used during the reverse process, allowing you to protect the source replica virtual server on the target back to the original source ESX host. If necessary, click **Provide Credentials** and specify a valid user on the virtual recovery appliance you have selected. The user must be a member of the Double-Take Admin security group. You can skip this step, however, you must have a reverse appliance selected before failover.
- **Send data to the source route when reversing the protection**—Specify an IP address on the source to route Double-Take data. 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.

Failover Monitor

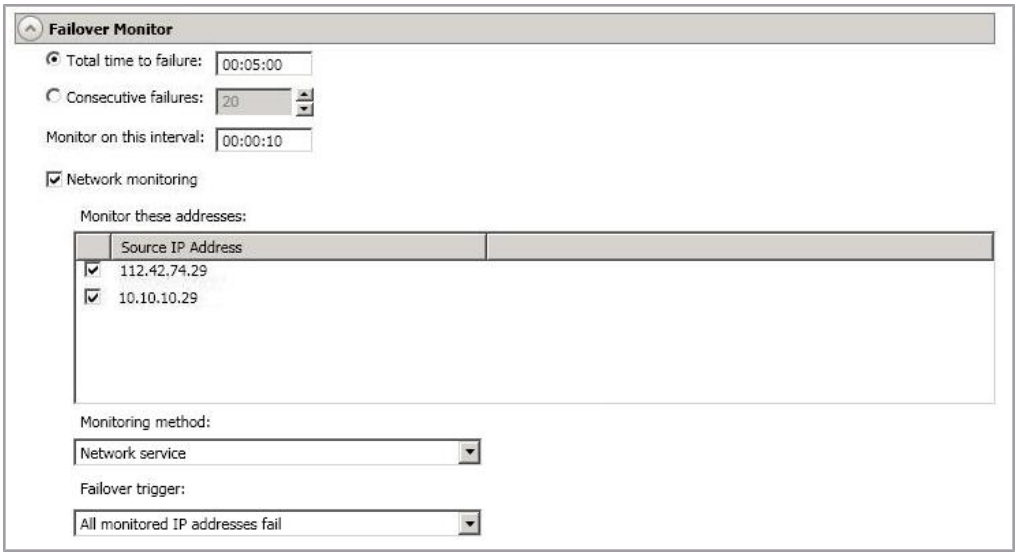

- **Total time to failure**—Specify, in hours:minutes:seconds, how long the target will keep trying to contact the source before the source is considered failed. This time is precise. If the total time has expired without a successful response from the source, this will be considered a failure.

  Consider a shorter amount of time for servers, such as a web server or order processing database, which must remain available and responsive at all times. Shorter times should be used where redundant interfaces and high-speed, reliable network links are available to prevent the false detection of failure. If the hardware does not support reliable communications, shorter times can lead to premature failover. Consider a longer amount of time for machines on slower networks or on a server that is not transaction critical. For example, failover would not be necessary in the case of a server restart.

- **Consecutive failures**—Specify how many attempts the target will make to contact the source before the source is considered failed. For example, if you have this option set to 20, and your source fails to respond to the target 20 times in a row , this will be considered a failure.
- **Monitor on this interval**—Specify, in hours:minutes:seconds, how long to wait between attempts to contact the source to confirm it is online. This means that after a response (success or failure) is received from the source, Double-Take will wait the specified interval time before contacting the source again. If you set the interval to 00:00:00, then a new check will be initiated immediately after the response is received.

  If you choose **Total time to failure**, do not specify a longer interval than failure time or your server will be considered failed during the interval period.

  If you choose **Consecutive failures**, your failure time is calculated by the length of time it takes your source to respond plus the interval time between each response, times the number of consecutive failures that can be allowed. That would be (response time + interval) * failure number. Keep in mind that timeouts from a failed check are included in the response time, so your failure time will not be precise.

- **Network monitoring**—With this option, the target will monitor the source using a network ping.
  - **Monitor these addresses**—Select each **Source IP Address** that you want the target to monitor. If you want to monitor additional addresses, enter the address and click **Add**.
  - **Monitoring method**—This option determines the type of network ping used for failover monitoring.
    - **Network service**—Source availability will be tested by an ICMP ping to confirm the route is active.
    - **Replication service**—Source availability will be tested by a UDP ping to confirm the Double-Take service is active.
    - **Network and replication services**—Source availability will be tested by both an ICMP ping to confirm the route is active and a UDP ping to confirm the Double-Take service is active. Both pings must fail in order to trigger a failover.
  - **Failover trigger**—If you are monitoring multiple IP addresses, specify when you want a failover condition to be triggered.
    - **One monitored IP address fails**—A failover condition will be triggered when any one of the monitored IP addresses fails. If each IP address is on a different subnet, you may want to trigger failover after one fails.
    - **All monitored IP addresses fail**—A failover condition will be triggered when all monitored IP addresses fail. If there are multiple, redundant paths to a server, losing one probably means an isolated network problem and you should wait for all IP addresses to fail.

Failover Options

Wait for user to initiate failover—By default, the failover process will wait for you to initiate it, allowing you to control when failover occurs. When a failure occurs, the job will wait in **Failover Condition Met** for you to manually initiate the failover process. Disable this option only if you want failover to occur immediately when a failure occurs.

Failover Identity

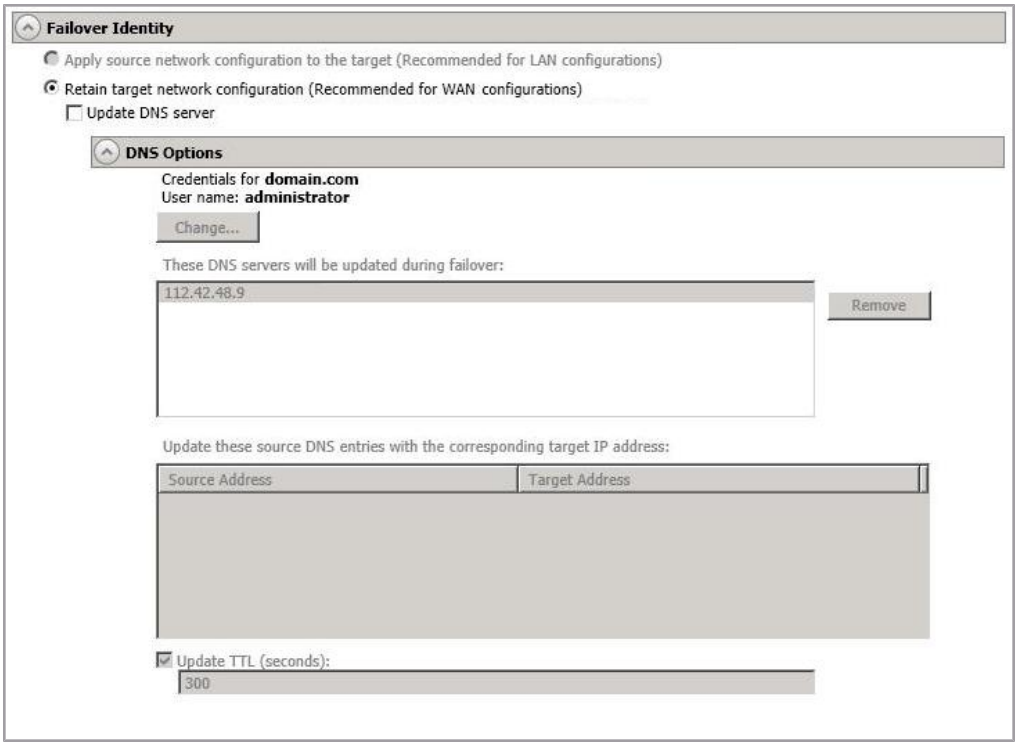

- **Retain target network configuration**—
  - **Update DNS server**—Specify if you want Double-Take to update your DNS server on . If DNS updates are made, the DNS records will be locked during . Be sure and review the *Core Double-Take requirements* on page 23 for the requirements for updating DNS.

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

    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.

If you select **Retain your target network configuration** but do not enable **Update DNS server**, you will need to specify failover scripts that update your DNS server during failover, or you can update the DNS server manually after failover. This would also apply to non--Microsoft Active Directory Integrated DNS servers. You will want to keep your target network configuration but do not update DNS. In this case, you will need to specify failover scripts that update your DNS server during failover, or you can update the DNS server manually after failover.

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

If you are using domain credentials during job creation, you must be able to resolve the domain name from the replica virtual machine using DNS before you can reverse.

Mirror, Verify & Orphaned Files

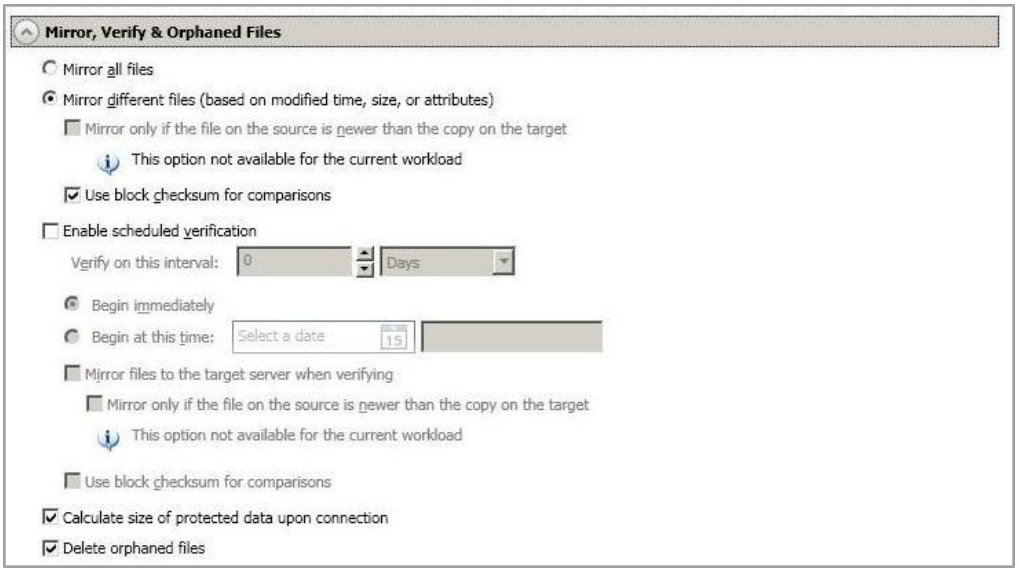

- **Mirror all files**—All protected files will be mirrored from the source to the target.
- **Mirror different files**—Only those protected files that are different based on date and time, size, or attributes will be mirrored from the source to the target.
  - **Mirror only if the file on the source is newer than the copy on the target**—This option is not available for V to ESX jobs.
  - **Use block checksum for comparisons**—For those files flagged as different, the mirroring process can perform a block checksum comparison and send only those blocks that are different.

File Differences Mirror Options Compared

The following table will help you understand how the various difference mirror options work together, including when you are using the block checksum option configured through the *Source server properties* on page 92.

An X in the table indicates that option is enabled. An X enclosed in parentheses (X) indicates that the option can be on or off without impacting the action performed during the mirror.

Not all job types have the source newer option available.

Source Server Properties	Job Properties			Action Performed
Block Checksum Option	File Differences Option	Source Newer Option	Block Checksum Option	
(X)	X			Any file that is different on the source and target based on the date, time, size, and/or attribute is transmitted to the target. The mirror sends the entire file.
(X)	X	X		Any file that is newer on the source than on the target based on date and/or time is transmitted to the target. The mirror sends the entire file.
	X		X	Any file that is different on the source and target based on date, time, size, and/or attributed is flagged as different. The mirror then performs a checksum comparison on the flagged files and only sends those blocks that are different.
X	X		X	The mirror performs a checksum comparison on all files and only sends those blocks that are different.
(X)	X	X	X	Any file that is newer on the source than on the target based on date and/or time is flagged as different. The mirror then performs a checksum comparison on the flagged files and only sends those blocks that are different.

- **Enable scheduled verification**—Verification is the process of confirming that the source replica data on the target is identical to the original 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, can automatically initiate a remirror, if configured. The remirror ensures data integrity between the source and target. When this option is enabled, Double-Take will verify the source replica data on the target and generate a verification log.

  

  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.

  - **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 at the specified date and time.
  - **Mirror files to the target server when verifying**—When this option is enabled, in addition to verifying the data and generating a log, Double-Take will also mirror to the target any protected files that are different on the source.
    - **Mirror only if the file on the source is newer than the copy on the target**—This option is not available for V to ESX jobs.
    - **Use block checksum for comparisons**—For those files flagged as different, the mirroring process can perform a block checksum comparison and send only those blocks that are different.
- **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 the calculation is enabled, it is completed before the job starts mirroring, which can take a significant amount of time depending on the number of files and system performance. If your job contains a large number of files, for example, 250,000 or more, you may want to disable the calculation so that data will start being mirrored sooner. 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 V to ESX jobs. The calculation time is when the system state protection processes hard links. If you disable the calculation, the hard link processing will not occur and you may have problems after failover, especially if your source is Windows 2008.

- **Delete orphaned files**—An orphaned file is a file that exists in the replica data on the

target, but does not exist in the protected data on the source. This option specifies if orphaned files should be deleted on the target during a mirror, verification, or restoration.

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

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 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.

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 95 for more information.

During a mirror, orphaned file processing success messages will be logged to a separate orphaned file log. 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 100 for details on the location of that folder. The orphaned log file is overwritten during each orphaned file processing during a mirror, and the log file will be a maximum of 50 MB.

Network Route

For **Send data to the target server using this route**, Double-Take will select, by default, a target route for transmissions. If desired, specify an alternate route on the target that the data will be transmitted through. This allows you to select a different route for Double-Take traffic. For example, you can separate regular network traffic and Double-Take traffic on a machine with multiple IP addresses.

The IP address used on the source will be determined through the Windows route table.

Compression

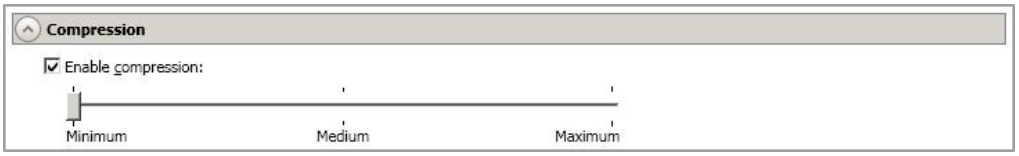

To help reduce the amount of bandwidth needed to transmit Double-Take data, compression allows you to compress data prior to transmitting it across the network. In a WAN environment this provides optimal use of your network resources. If compression is enabled, the data is compressed before it is transmitted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

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

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

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

Bandwidth

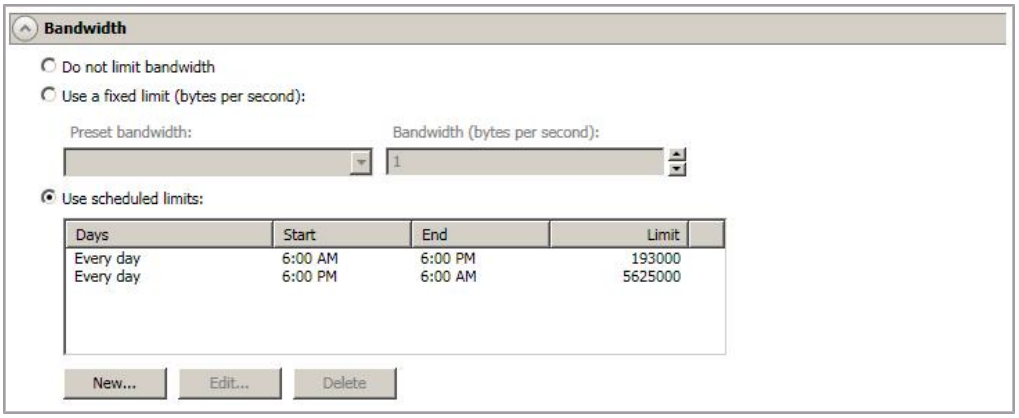

Bandwidth limitations are available to restrict the amount of network bandwidth used for Double-Take data transmissions. When a bandwidth limit is specified, Double-Take never exceeds that allotted amount. The bandwidth not in use by Double-Take is available for all other network traffic.

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

- **Do not limit bandwidth**—Double-Take will transmit data using 100% bandwidth availability.
- **Use a fixed limit**—Double-Take will transmit data using a limited, fixed bandwidth. Select a **Preset bandwidth** limit rate from the common bandwidth limit values. The **Bandwidth** field will automatically update to the bytes per second value for your selected bandwidth. This is the maximum amount of data that will be transmitted per second. If desired, modify the bandwidth using a bytes per second value. The minimum limit should be 3500 bytes per second.
- **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**.

4. If you want to modify the workload items or replication rules for the job, click **Edit workload or replication rules**. Modify the **Workload item** you are protecting, if desired. Additionally, you can modify the specific **Replication Rules** for your job.

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

   You can also enter wildcard rules, however you should do so carefully. Rules are applied to files that are closest in the directory tree to them. If you have rules that include multiple folders, an exclusion rule with a wild card will need to be added for each folder that it needs applied to. For example, if you want to exclude all .log files from D:\ and your rules include D:\, D:\Dir1 , and D:\Dir2, you would need to add the exclusion rule for the root and each subfolder rule. So you will need to add exclude rules for D:*.log , D:\Dir1*.log, and D:\Dir2*.log.

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

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

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

   Errors are designated by a white X inside a red circle. Warnings are designated by a black exclamation point (!) inside a yellow triangle. A successful validation is designated by a white checkmark inside a green circle. You can sort the list by the icon to see errors, warnings, or successful validations together. Click on any of the validation items to see details. You must correct any errors before you can enable protection.

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

Viewing a V to ESX job log

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

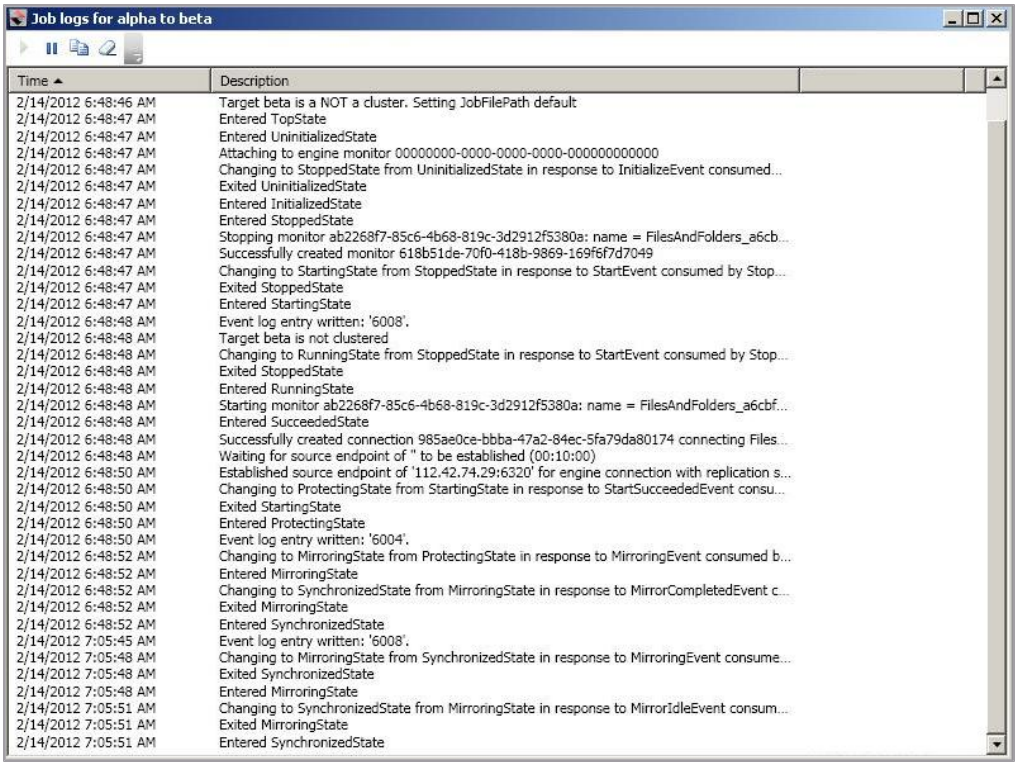

The following table identifies the controls and the table columns in the **Job logs** window.

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.

Failing over V to ESX jobs

When a failover condition has been met, failover will be triggered automatically if you disabled the wait for user option during your failover configuration. If the wait for user before failover option is enabled, you will be notified in the console when a failover condition has been met. At that time, you will need to trigger it manually from the console when you are ready.

1. On the **Manage Jobs** page, highlight the job that you want to failover and click **Failover or Cutover** in the toolbar.
2. Select the type of failover to perform.
   - **Failover to live data**—Select this option to initiate a full, live failover using the current data on the target. This option will shutdown the source machine (if it is online), stop the protection job, and start the replica virtual machine on the target with full network connectivity.
   - **Perform test failover**—Select this option to perform a test failover using the current data on the target. This option will leave the source machine online, stop the protection job, and start the replica virtual machine on the target without network connectivity.
   - **Failover to a snapshot**—This option is not available for V to ESX jobs.
3. Select how you want to handle the data in the target queue. You may want to check the amount of data in queue on the target by reviewing the *Statistics* on page 688 or *Performance Monitor* on page 790.
   - **Apply data in target queues before failover or cutover**—All of the data in the target queue will be applied before failover begins. The advantage to this option is that all of the data that the target has received will be applied before failover begins. 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 failover or cutover immediately**—All of the data in the target queue will be discarded and failover will begin immediately. The advantage to this option is that failover will occur immediately. The disadvantage is that any data in the target queue will be lost.
   - **Revert to last good snapshot if target data state is bad**—If the target data is in a bad state, Double-Take will automatically revert to the last good Double-Take snapshot before failover begins. If the target data is in a good state, Double-Take will not revert the target data. Instead, Double-Take will apply the data in the target queue and then failover. The advantage to this option is that good data on the target is guaranteed to be used. The disadvantage is that if the target data state is bad, you will lose any data between the last good snapshot and the failure.
4. When you are ready to begin failover, click **Failover**.

Depending on your replica configuration, you may have to reboot your replica after failover. You will be prompted to reboot if it is necessary.

If your source was using vCenter, you may have problems with failover if vCenter is down or if it is unreachable. See the technical support article 34768 for details on how to complete failover in this situation.

If you used domain credentials when you originally created your job, you must be able to resolve the domain name from the replica virtual machine using DNS before you can reverse.

5. If you performed a test failover, you can undo it by selecting **Undo Failover or Cutover** in the toolbar. The replica virtual machine on the target will be shut down and the protection job will be restarted performing a file differences mirror.

Reversing V to ESX jobs

Reversing protection allows you to protect your source replica virtual server running on the target back to the original source host. Your original source host must have a licensed virtual recovery appliance like your original target host. See *V to ESX requirements* on page 493.

1. Make sure you original source virtual machine is not running.
2. If you used domain credentials when you originally created your job, make sure you can resolve the domain name from the replica virtual machine using DNS.
3. On the **Manage Jobs** page, highlight the job that you want to reverse and click **Reverse** in the toolbar. The flow of mirroring and replication data will change. Data will be transmitted from the replica virtual machine on the target back to the original source host.

After the reverse is complete, your source replica virtual machine on the target is being protected to your original source host. In the event you want to go back to your original server roles and hardware configuration, you can failover again.

The original disk structure on the original source virtual machine will be deleted and re-created the first time you perform a reverse. This reverse will perform a full mirror. Subsequent reverses will always perform a file differences mirror.

Chapter 14 V to Hyper-V protection

Create a V to Hyper-V job when you want to protect a virtual machine on a Hyper-V host to another Hyper-V host. This protection type is at the virtual machine guest level. You can reverse the protection after failover has occurred.

- See *V to Hyper-V requirements* on page 545—V to Hyper-V protection includes specific requirements for this type of protection.
- See *Creating a V to Hyper-V job* on page 547—The section includes step-by-step instructions for creating a V to Hyper-V job.
- See *Managing and controlling V to Hyper-V jobs* on page 563—You can view status information about your V to Hyper-V jobs and learn how to control these jobs.
- See *Failing over V to Hyper-V jobs* on page 600—Use this section when a failover condition has been met or if you want to failover manually.
- See *Reversing V to Hyper-V jobs* on page 602—Use this section to reverse protection. The source replica on the target is now sending data back to the original source.

V to Hyper-V requirements

In addition to the *Core Double-Take requirements* on page 23, use these requirements for V to Hyper-V protection.

- **Source server**—The source server can be any virtual server running any of the operating systems listed in the *Core Double-Take requirements* on page 23. However, if you are using a Windows 2003 operating system, you must have Service Pack 2 which is required for Hyper-V Integration Services. Additionally, if your source is a virtual server and you want Double-Take to monitor it for failover, then you must have Integration Components installed on the guest operating system and the virtual machine must be powered on.
- **Target server**—The target server can be any Windows 2008, 2008 R2, 2012, or 2012 R2 operating system from the *Core Double-Take requirements* on page 23 that has the Hyper-V role enabled. In addition, you can use Hyper-V Server 2008 R2, Server Core 2008 R2, Server Core 2012, or Server Core 2012 R2 with the Hyper-V role enabled. (Hyper-V Server 2008 and Server Core 2008 are not supported.) In each case, the source and target must be running identical operating system versions. For example, your source cannot be Windows 2008 (or Windows 2012) and your target Windows 2008 R2 (or Windows 2012 R2).
- **Server Core**—In addition to the Server Core requirements above, there is a Server Core limitation. DNS updates are not supported for Server Core servers.
- **Disk types**—Virtual machines can use raw, pass-through, or differencing disks, however, they will be virtual hard disks on the replica on the target.
- **IP addressing**—IPv4 is the only supported IP version.
- **Microsoft .NET Framework**—Microsoft .NET Framework version 3.5 Service Pack 1 is required. This version is not included in the .NET version 4.0 release. Therefore, even if you have .NET version 4.0 installed, you will also need version 3.5.1. For Windows 2008 and earlier, you can install this version from the Double-Take DVD, via a web connection during the Double-Take installation, or from a copy you have obtained manually from the Microsoft web site. For Windows 2008 R2 and later, you need to enable it through Windows features.
- **Live migration**—Windows 2008 R2 and 2012 live migration is supported for CSV clustered virtual machines, but the shared-nothing configuration is not supported.
- **Snapshots**—Double-Take snapshots are not supported with V to Hyper-V jobs.
- **Supported configurations**—The following table identifies the supported configurations for a V to Hyper-V job. (These are source and target configurations for the host, not the virtual machines.)

Configuration		Supported	Not Supported
Source to target configuration[1]	One to one, active/standby	X	
	One to one, active/active		X
	Many to one	X	
	One to many	X	
	Chained		X
	Single server	X	
Server configuration	Standalone to standalone	X	
	Standalone to cluster	X	
	Cluster to standalone	X	
	Cluster to cluster	X	
	Cluster Shared Volumes (CSV) guest level	X	
	Cluster Shared Volumes (CSV) host level		X
Upgrade configuration	Upgrade 6.0 V to Hyper-V job to 7.0 V to Hyper-V job	X	
Version 7.0 console[2,3]	Version 7.0 console can create job for 5.3 source and 5.3 target		X
	Version 7.0 console can create job for 6.0 source and 6.0 target	X	
	Version 7.0 console can create job for 7.0 source and 7.0 target	X	

1. See *Supported configurations* on page 16 for details on each of the source to target configurations.
2. Once you upgrade your console to version 7.0, existing jobs that are running version 5.3 will not appear in the console until the target of the job is upgraded to version 7.0.
3. Newer job options available in the version 7.0 console will not be functional when creating jobs for servers running version 6.0.

Creating a V to Hyper-V job

Use these instructions to create a V to Hyper-V job.

1. Click **Get Started** from the toolbar.
2. Select **Double-Take Availability** and click **Next**.
3. Select **Protect files and folders, an application, or an entire Windows server** and click **Next**.
4. Choose your source server. This is the Hyper-V server or cluster that is hosting the virtual machines that you want to protect.

- **Current Servers**—This list contains the servers currently available in your console session. Servers that are not licensed for the workflow you have selected will be filtered out of the list. Select your source server from the list.
- **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 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.

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. The user must also have administrative rights for Microsoft Hyper-V.

5. Click **Next** to continue.
6. Choose the type of workload that you want to protect. Under **Server Workloads**, in the **Workload types** pane, select **Guest-Level Hyper-V**. In the **Workload items** pane, select the virtual machines on the Hyper-V source that you want to protect. The list of virtual machines will vary depending on whether your source is a Hyper-V server, cluster, or node.

   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.

7. Click **Next** to continue.
8. Choose your target server. This is the Hyper-V server that will store the replicas of the virtual machines 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.
- **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. The user must also have administrative rights for Microsoft Hyper-V.

---

9. Click **Next** to continue.

10. You have many options available for your V to Hyper-V jobs. 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 561.

    - *Replica Virtual Machines* on page 551
    - *Replica Virtual Machine Configuration* on page 552
    - *Mirror, Verify & Orphaned Files* on page 553
    - *Network Route* on page 557
    - *Reverse Network Route* on page 558
    - *Compression* on page 559
    - *Bandwidth* on page 560

Replica Virtual Machines

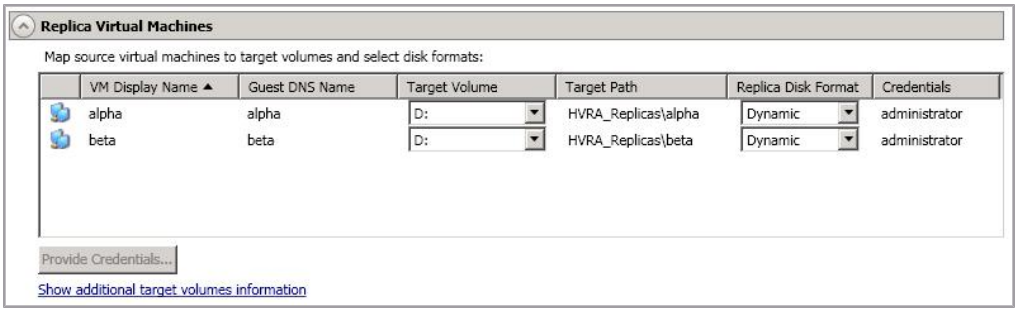

- **Target Volume**—For each virtual machine you are protecting, specify the volume where you want to store the replica virtual machine on the target.
- **Target Path**—For each virtual machine you are protecting, specify a path on the selected **Target Volume** where you want to store the replica virtual machine on the target.
- **Replica Disk Format**—For each virtual machine you are protecting, specify the type of disk, **Dynamic** or **Fixed**, that will be created on the replica virtual machine.
- **Provide Credentials**—If necessary, click **Provide Credentials** and specify a valid user on the virtual recovery appliance you have selected. The user must be a member of the Double-Take Admin security group and must have administrative rights for Hyper-V.
- **Show additional target volumes information**—Click this link to see storage information for the volumes on your target. This will help you select the appropriate volumes for your replica virtual machines.

If your target is a cluster, you will only see the clustered volumes that have been added to the cluster prior to creating the job. If the target is a cluster node, you will only see non-clustered volumes.

Replica Virtual Machine Configuration

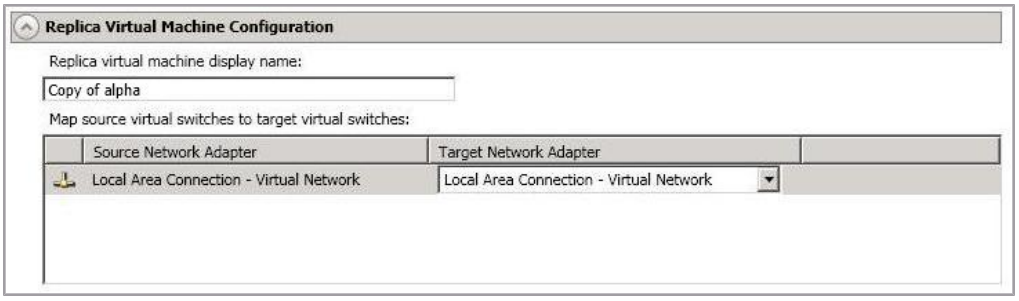

- **Replica virtual machine display name**—Specify the name of the replica virtual machine. This will be the display name of the virtual machine on the host system.
- **Map source virtual switches to target virtual switches**—Identify how you want to handle the network mapping after failover. The **Source Network Adapter** column lists the NICs from the source. Map each one to a **Target Network Adapter**, which is a virtual network on the target. You can also choose to discard the source's NIC and IP addresses, or you can to failover the NIC and IP addresses but leave them in a not connected state.

Mirror, Verify & Orphaned Files

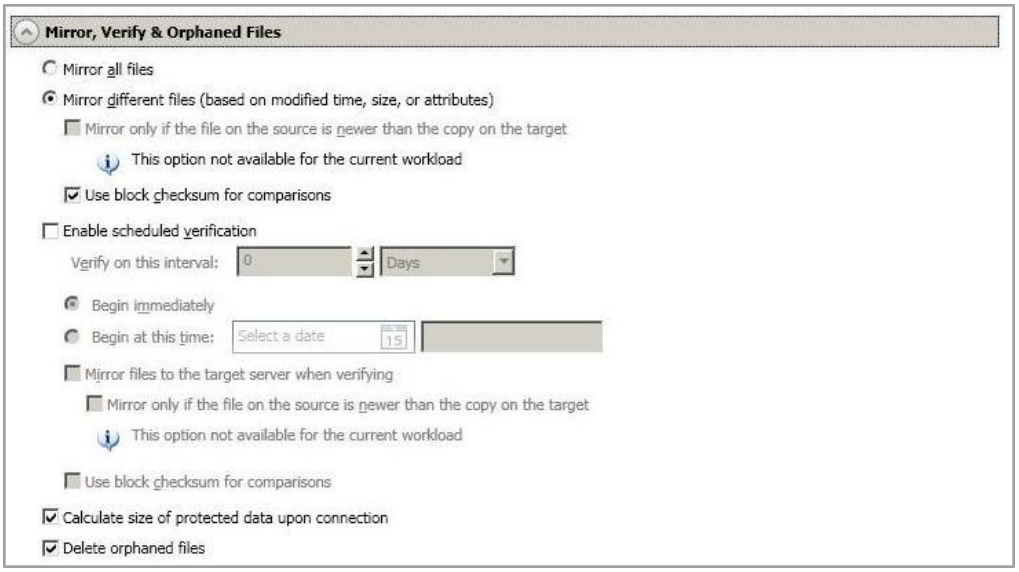

- **Mirror all files**—All protected files will be mirrored from the source to the target.
- **Mirror different files**—Only those protected files that are different based on date and time, size, or attributes will be mirrored from the source to the target.
  - **Mirror only if the file on the source is newer than the copy on the target**—This option is not available for V to Hyper-V jobs.
  - **Use block checksum for comparisons**—For those files flagged as different, the mirroring process can perform a block checksum comparison and send only those blocks that are different.

File Differences Mirror Options Compared

The following table will help you understand how the various difference mirror options work together, including when you are using the block checksum option configured through the *Source server properties* on page 92.

An X in the table indicates that option is enabled. An X enclosed in parentheses (X) indicates that the option can be on or off without impacting the action performed during the mirror.

Not all job types have the source newer option available.

Source Server Properties	Job Properties			Action Performed
Block Checksum Option	File Differences Option	Source Newer Option	Block Checksum Option	
(X)	X			Any file that is different on the source and target based on the date, time, size, and/or attribute is transmitted to the target. The mirror sends the entire file.
(X)	X	X		Any file that is newer on the source than on the target based on date and/or time is transmitted to the target. The mirror sends the entire file.
	X		X	Any file that is different on the source and target based on date, time, size, and/or attributed is flagged as different. The mirror then performs a checksum comparison on the flagged files and only sends those blocks that are different.
X	X		X	The mirror performs a checksum comparison on all files and only sends those blocks that are different.
(X)	X	X	X	Any file that is newer on the source than on the target based on date and/or time is flagged as different. The mirror then performs a checksum comparison on the flagged files and only sends those blocks that are different.

- **Enable scheduled verification**—Verification is the process of confirming that the source replica data on the target is identical to the original 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, can automatically initiate a remirror, if configured. The remirror ensures data integrity between the source and target. When this option is enabled, Double-Take will verify the source replica data on the target and generate a verification log.

  

  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.

  - **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 at the specified date and time.
  - **Mirror files to the target server when verifying**—When this option is enabled, in addition to verifying the data and generating a log, Double-Take will also mirror to the target any protected files that are different on the source.
    - **Mirror only if the file on the source is newer than the copy on the target**—This option is not available for V to Hyper-V jobs.
    - **Use block checksum for comparisons**—For those files flagged as different, the mirroring process can perform a block checksum comparison and send only those blocks that are different.
- **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 the calculation is enabled, it is completed before the job starts mirroring, which can take a significant amount of time depending on the number of files and system performance. If your job contains a large number of files, for example, 250,000 or more, you may want to disable the calculation so that data will start being mirrored sooner. 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 V to Hyper-V jobs. The calculation time is when the system state protection processes hard links. If you disable the calculation, the hard link processing will not occur and you may have problems after failover, especially if your source is Windows 2008.

- **Delete orphaned files**—An orphaned file is a file that exists in the replica data on the

target, but does not exist in the protected data on the source. This option specifies if orphaned files should be deleted on the target during a mirror, verification, or restoration.

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

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 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.

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 95 for more information.

During a mirror, orphaned file processing success messages will be logged to a separate orphaned file log. 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 100 for details on the location of that folder. The orphaned log file is overwritten during each orphaned file processing during a mirror, and the log file will be a maximum of 50 MB.

Network Route

This section is not applicable if your target is a cluster.

For **Send data to the target server using this route**, Double-Take will select, by default, a target route for transmissions. If desired, specify an alternate route on the target that the data will be transmitted through. This allows you to select a different route for Double-Take traffic. For example, you can separate regular network traffic and Double-Take traffic on a machine with multiple IP addresses.

The IP address used on the source will be determined through the Windows route table.

Reverse Network Route

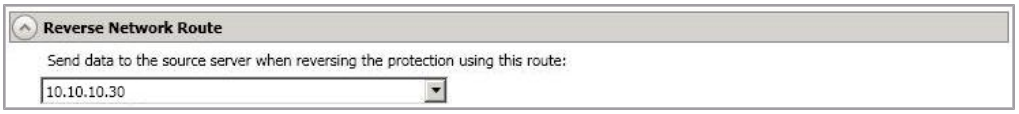

This section is not applicable if your source is a cluster.

Send data to the source server when reversing the protection using this route—By default, Double-Take will select a default source route for transmissions. If desired, specify an alternate route on the source 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.

The IP address used on the target during a reverse will be determined through the Windows route table.

Compression

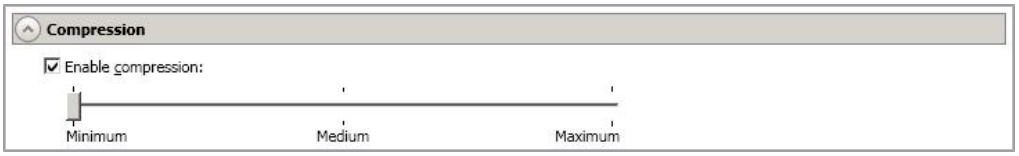

To help reduce the amount of bandwidth needed to transmit Double-Take data, compression allows you to compress data prior to transmitting it across the network. In a WAN environment this provides optimal use of your network resources. If compression is enabled, the data is compressed before it is transmitted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

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

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

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

Bandwidth

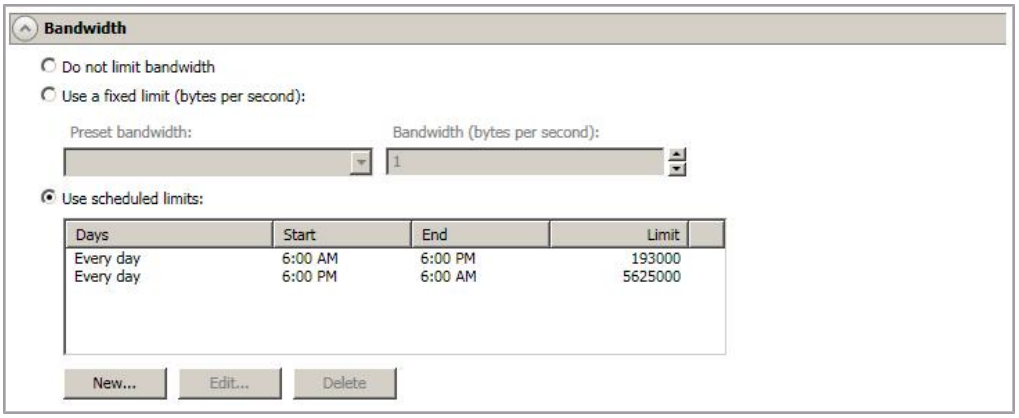

Bandwidth limitations are available to restrict the amount of network bandwidth used for Double-Take data transmissions. When a bandwidth limit is specified, Double-Take never exceeds that allotted amount. The bandwidth not in use by Double-Take is available for all other network traffic.

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

- **Do not limit bandwidth**—Double-Take will transmit data using 100% bandwidth availability.
- **Use a fixed limit**—Double-Take will transmit data using a limited, fixed bandwidth. Select a **Preset bandwidth** limit rate from the common bandwidth limit values. The **Bandwidth** field will automatically update to the bytes per second value for your selected bandwidth. This is the maximum amount of data that will be transmitted per second. If desired, modify the bandwidth using a bytes per second value. The minimum limit should be 3500 bytes per second.
- **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**.

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

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

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

Because this job type allowed you to create more than one job at a time, options that are job specific were not presented during the job creation process. Once the jobs are created, you can

edit each one individually to see and edit all of the various job options. See *Editing a V to Hyper-V job* on page 577.

Managing and controlling V to Hyper-V jobs

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

The jobs displayed in the right pane depend on the server group folder selected in the left pane. Every job for each server in your console session is displayed when the **Jobs on All Servers** group is selected. If you have created and populated server groups (see *Managing servers* on page 65), then only the jobs associated with the server or target servers in that server group will be displayed in the right pane.

- See *Overview job information displayed in the top pane* on page 563
- See *Detailed job information displayed in the bottom pane* on page 565
- See *Job controls* on page 567

Overview job information displayed in the top pane

The top pane displays high-level overview information about your jobs.

Column 1 (Blank)

The first blank column indicates the state of the job.

 The job is in a healthy state.

 The job is in a warning state. This icon is also displayed on any server groups that you have created that contain a job in a warning state.

 The job is in an error state. This icon is also displayed on any server groups that you have created that contain a job in an error state.

 The job is in an unknown state.

Job

The name of the job

Source Server

The name of the source.

Target Server

The name of the target.

Job Type

Each job type has a unique job type name. This job is a V to Hyper-V job. For a complete list of all job type names, press F1 to view the Double-Take Console online help.

Activity

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

Mirror Status

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

Replication Status

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

Transmit Mode

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

Detailed job information displayed in the bottom pane

The details displayed in the bottom pane of the **Manage Jobs** page provide additional information for the job highlighted in the top pane. If you select multiple jobs, the details for the first selected job will be displayed.

Name

The name of the job

Target data state

- **OK**—The data on the target is in a good state.
- **Mirroring**—The target is in the middle of a mirror process. The data will not be in a good state until the mirror is complete.
- **Mirror Required**—The data on the target is not in a good state because a remirror is required. This may be caused by an incomplete or stopped mirror or an operation may have been dropped on the target.
- **Restore Required**—The data on the source and target do not match because of a failover condition. Restore the data from the target back to the source. If you want to discard the changes on the target, you can remirror to resynchronize the source and target.
- **Snapshot Reverted**—The data on the source and target do not match because a snapshot has been applied on the target. Restore the data from the target back to the source. If you want to discard the changes on the target, you can remirror to resynchronize the source and target.
- **Busy**—The source is low on memory causing a delay in getting the state of the data on the target.
- **Not Loaded**—Double-Take target functionality is not loaded on the target server. This may be caused by an activation code error.
- **Unknown**—The console cannot determine the status.

Mirror remaining

The total number of mirror bytes that are remaining to be sent from the source to the target

Mirror skipped

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

Replication queue

The total number of replication bytes in the source queue

Disk queue

The amount of disk space being used to 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**.

Connected since

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

Recent activity

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

Additional information

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

Job controls

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

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

Create a New Job 

This button leaves the **Manage Jobs** page and opens the **Get Started** page.

View Job Details 

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

Delete 

Stops (if running) and deletes the selected jobs.

If you no longer want to protect the source and no longer need the replica of the source on the target, select to delete the associated replica virtual machine. Selecting this option will remove the job and completely delete the replica virtual machine on the target.

If you no longer want to mirror and replicate data from the source to the target but still want to keep the replica of the source on the target, select to keep the associated replica virtual machine. You may want to use this option to relocate the virtual hard disks and create a new job between the original source and the new location. Selecting this option, will preserve the source replica on the target.

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 77. You will remain on the **Manage Jobs** page after updating the server credentials. If your servers use the same credentials, make sure you also update the credentials on the **Manage Servers** page so that the Double-Take Console can authenticate to the servers in the console session. See *Managing servers* on page 65.

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. Failover monitoring will continue 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. Failover monitoring will continue while the job is stopped.

Take Snapshot 

Snapshots are not applicable to V to Hyper-V jobs.

Manage Snapshots 

Snapshots are not applicable to V to Hyper-V jobs.

Failover or Cutover 

Starts the failover process. See *Failing over V to Hyper-V jobs* on page 600 for the process and details of failing over a V to Hyper-V job.

Failback 

Starts the failback process. Failback does not apply to V to Hyper-V jobs.

Restore 

Starts the restoration process. Restore does not apply to V to Hyper-V jobs.

Reverse 

Reverses protection. The job will start mirroring in the reverse direction with the job name and log file names changing accordingly. After the mirror is complete, the job will continue running in the opposite direction. See *Reversing V to Hyper-V jobs* on page 602 for the process and details of reversing a V to Hyper-V job.

Undo Failover 

Cancels a test failover by undoing it. This resets the servers and the job back to their original state.. See *Failing over V to Hyper-V jobs* on page 600 for the process and details of undoing a failed over V to Hyper-V job.

View Job Log 

Opens the job log. On the right-click menu, this option is called **View Logs**, and you have the option of opening the job log, source server log, or target server log. See *Viewing the log files through the Double-Take Console* on page 678 for details on all three of these logs.

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 all files**—All protected files will be mirrored from the source to the target.
  - **Mirror different files**—Only those protected files that are different based on date and time, size, or attributes will be mirrored from the source to the target.
    - **Mirror only if the file on the source is newer than the copy on the target**—This option is available for V to Hyper-V jobs, but ideally it should not be used.

- **Use block checksum for comparisons**—For those files flagged as different, the mirroring process can perform a block checksum comparison and send only those blocks that are different.

- **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 the calculation is enabled, it is completed before the job starts mirroring, which can take a significant amount of time depending on the number of files and system performance. If your job contains a large number of files, for example, 250,000 or more, you may want to disable the calculation so that data will start being mirrored sooner. 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 V to Hyper-V jobs. The calculation time is when the system state protection processes hard links. If you disable the calculation, the hard link processing will not occur and you may have problems after failover, especially if your source is Windows 2008.

- **Verify**—Even if you have scheduled the verification process, you can run it manually any time a mirror is not in progress.
  - **Create verification report only**—This option verifies the data and generates a verification log, but it does not remirror any files that are different on the source and target. See *Verification log* on page 102 for details on the log file.
  - **Mirror files to the target server automatically**—When this option is enabled, in addition to verifying the data and generating a log, Double-Take will also mirror to the target any protected files that are different on the source.
    - **Mirror only if the file on the source is newer than the copy on the target**—This option is available for V to Hyper-V jobs, but ideally it should not be used.
    - **Use block checksum for comparisons**—For those files flagged as different, the mirroring process can perform a block checksum comparison and send only those blocks that are different.
- **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.
- **Update Shares**—Shares are not applicable because they are automatically included with the system state that is being protected with the entire server.

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 65.

Type a server name

Displays only jobs that contain the text you entered. If you have created and populated server groups, then only jobs that contain the text you entered associated with the server or target servers in that server group will be displayed. See *Managing servers* on page 65.

Overflow Chevron 

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

Viewing V to Hyper-V job details

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

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

Job name

The name of the job

Job type

Each job type has a unique job type name. This job is a V to Hyper-V job. For a complete list of all job type names, press F1 to view the Double-Take Console online help.

Health

 The job is in a healthy state.

 The job is in a warning state.

 The job is in an error state.

 The job is in an unknown state.

Activity

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

Connection ID

The incremental counter used to number connections. The number is incremented when a connection is created. It is also incremented by internal actions, such as an auto-disconnect and auto-reconnect. The lowest available number (as connections are created, stopped, deleted, and so on) will always be used. The counter is reset to one each time the Double-Take service is restarted.

Transmit mode

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

Target data state

- **OK**—The data on the target is in a good state.
- **Mirroring**—The target is in the middle of a mirror process. The data will not be in a good state until the mirror is complete.
- **Mirror Required**—The data on the target is not in a good state because a remirror is required. This may be caused by an incomplete or stopped mirror or an operation may have been dropped on the target.
- **Restore Required**—The data on the source and target do not match because of a failover condition. Restore the data from the target back to the source. If you want to discard the changes on the target, you can remirror to resynchronize the source and target.
- **Snapshot Reverted**—The data on the source and target do not match because a snapshot has been applied on the target. Restore the data from the target back to the source. If you want to discard the changes on the target, you can remirror to resynchronize the source and target.
- **Busy**—The source is low on memory causing a delay in getting the state of the data on the target.
- **Not Loaded**—Double-Take target functionality is not loaded on the target server. This may be caused by an activation code error.
- **Unknown**—The console cannot determine the status.

Target route

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

Compression

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

Encryption

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

Bandwidth limit

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

Connected since

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

Additional information

Depending on the current state of your job, you may see additional information displayed to keep you informed about the progress and status of your job. If there is no

additional information, you will see (None) displayed.

Mirror status

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

Mirror percent complete

The percentage of the mirror that has been completed

Mirror remaining

The total number of mirror bytes that are remaining to be sent from the source to the target

Mirror skipped

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

Replication status

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

Replication queue

The total number of replication bytes in the source queue

Disk queue

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

Bytes sent

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

Bytes sent compressed

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

Validating a V to Hyper-V job

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

1. From the **Manage Jobs** page, highlight the job and click **View Job Details** in the toolbar.
2. In the **Tasks** area on the right on the **View Job Details** page, click **Validate job properties**.
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 enable protection. Depending on the error, you may be able to click **Fix** or **Fix All** and let Double-Take correct the problem for you. For those errors that Double-Take cannot correct automatically, you will need to modify the source or target to correct the error, or you can select a different target. You must revalidate the selected servers, by clicking **Recheck**, until the validation check passes without errors.

   Validation checks for an existing job are logged to the job log on the target server. See *Log files* on page 677 for details on the various log files.
4. Once your servers have passed validation, click **Close**.

Editing a V to Hyper-V job

Use these instructions to edit a V to Hyper-V job.

1. From the **Manage Jobs** page, highlight the job and click **View Job Details** in the toolbar.
2. In the **Tasks** area on the right on the **View Job Details** page, click **Edit job properties**. (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. Because you configured multiple jobs at once when you first established your protection, not all of the individual job options were available during job creation. Therefore, you will have additional job options when you edit an existing job.

   

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

   There will be additional sections on the **Edit Job Properties** page that you will be able to view only. You cannot edit those sections.

   Go to each page identified below to see the options that you can edit on the **Edit Job Properties** page. After you have configured your options, continue with the next step on page 597.

   - *General* on page 578
   - *Replica Virtual Machine Location* on page 579
   - *Replica Virtual Machine Configuration* on page 580
   - *Replica Virtual Machine Volumes* on page 581
   - *Replica Virtual Machine Network Settings* on page 582
   - *Failover Monitor* on page 583
   - *Failover Options* on page 585
   - *Failover Identity* on page 586
   - *Mirror, Verify & Orphaned Files* on page 588
   - *Network Route* on page 592
   - *Reverse Network Route* on page 593
   - *Compression* on page 594
   - *Bandwidth* on page 595

General

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

Replica Virtual Machine Location

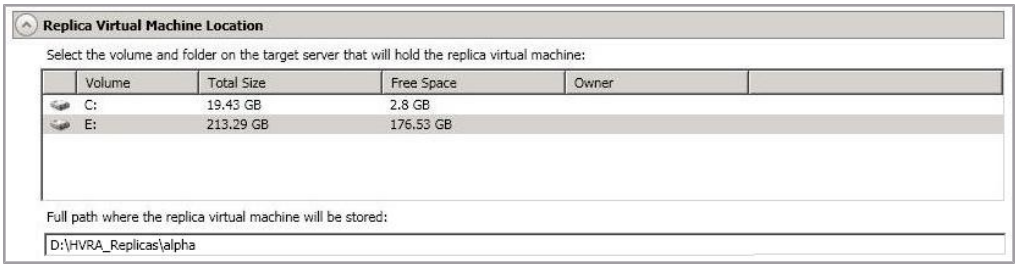

- **Select the volume and folder on the target server that will hold the replica virtual machine**—Select one of the volumes from the list to indicate the volume on the target where you want to store the new virtual server when it is created. The target volume must have enough **Free Space** to store the source data.
- **Full path where the replica virtual machine will be stored**—Specify a location on the selected **Volume** to store the replica of the source. By specifying an existing folder, you can reuse an existing virtual machine on your Hyper-V target created by a previous protection job. This can be useful for pre-staging data on a virtual machine over a LAN connection and then relocating it 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. In order to use a pre-existing virtual disk, it must be a valid virtual disk and it cannot be attached to any registered virtual machine. In a WAN environment, you may want to take advantage of re-using an existing virtual disk by using a process similar to the following.
    a. Create a protection job in a LAN environment, letting Double-Take create the virtual disk for you.
    b. Complete the mirror process locally.
    c. Delete the protection job and when prompted, select to keep the replica.
    d. From the Hyper-V Manager, delete the replica virtual machine, which will delete the virtual machine configuration but will keep the associated hard disk files.
    e. Shut down and move the Hyper-V target server to your remote site.
    f. After the Hyper-V target server is back online at the remote site, create a new protection job for the same source server. Double-Take will reuse the existing hard disk files and perform a difference mirror over the WAN to bring the virtual machine up-to-date.

Replica Virtual Machine Configuration

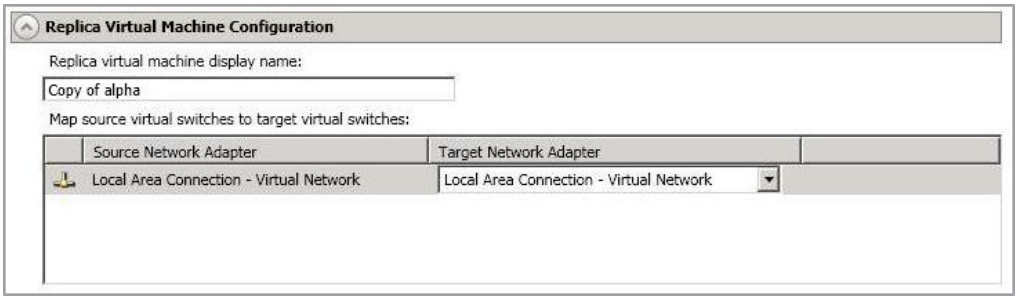

- **Replica virtual machine display name**—Specify the name of the replica virtual machine. This will be the display name of the virtual machine on the host system.
- **Map source virtual switches to target virtual switches**—Identify how you want to handle the network mapping after failover. The **Source Network Adapter** column lists the NICs from the source. Map each one to a **Target Network Adapter**, which is a virtual network on the target. You can also choose to discard the source's NIC and IP addresses, or you can to failover the NIC and IP addresses but leave them in a not connected state.

Replica Virtual Machine Volumes

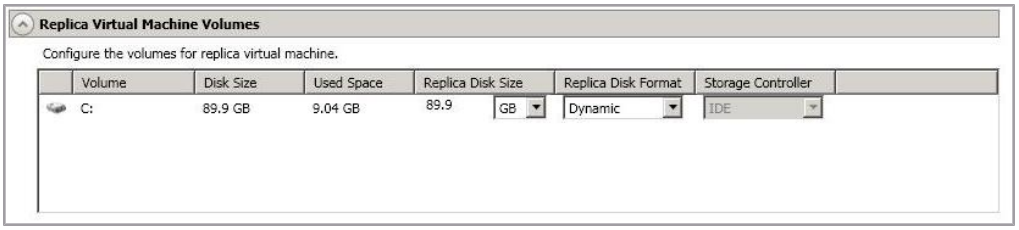

- **Replica Disk Size**—For each volume you are protecting, specify the size of the replica disk on the target. Be sure and include the value in MB or GB for the disk. The value must be at least the size of the specified **Used Space** on that volume.

In some cases, the replica virtual machine may use more virtual disk space than the size of the source volume due to differences in how the virtual disk's block size is formatted and how hard links are handled. To avoid this issue, specify the size of your replica to be at least 5 GB larger.

Snapshots are stored on the replica, so if you enable snapshots, be sure that you configure your replica virtual machine disk size large enough to maintain snapshots.

- **Replica Disk Format**—For each volume you are protecting, specify the format of the disk, **Dynamic** or **Fixed**, that will be created on the replica virtual machine. Any disk format specification will be discarded if you are reusing a disk from the **Full path where the replica virtual machine will be stored** from the **Replica Virtual Machine Location** section.
- **Storage Controller**—For each volume you are protecting, specify the type of **Storage Controller** that you want to use for each volume on the target. If your virtual machine is a Generation 2 VM (Windows 2012 or later), SCSI is the only controller option.

The system volume must be an IDE controller. In addition, up to two more volumes can be attached to an IDE controller. If you are protecting more than three volumes on the source, you will need to install the Hyper-V Integration Components to acquire a SCSI device after failover to attach these volumes to the replica virtual machine. You must be using Windows 2003 Service Pack 2 or later to use Hyper-V Integration Components.See your Microsoft documentation for more information.

Replica Virtual Machine Network Settings

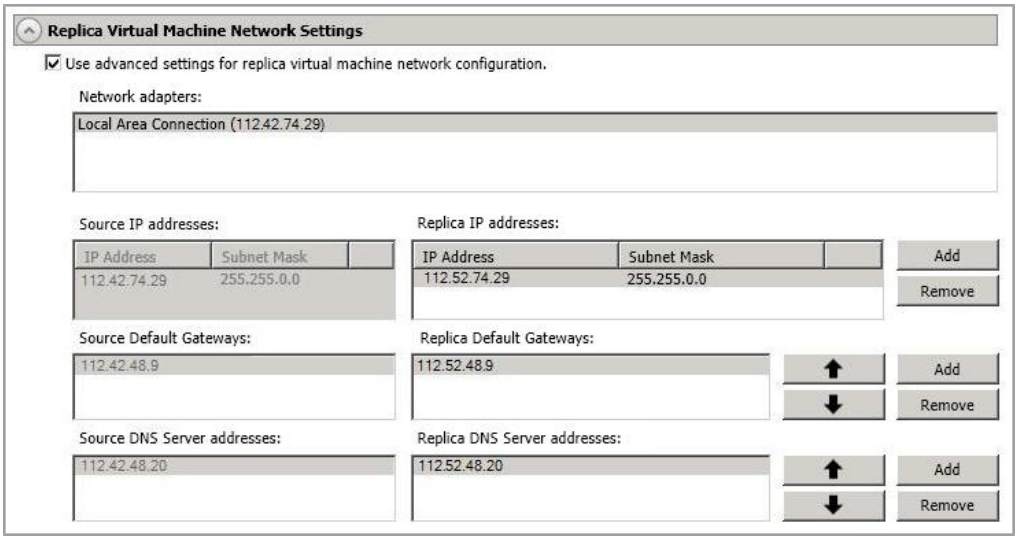

- **Use advanced settings for replica virtual machine network configuration**—Select this option to enable the replica virtual machine network setting configuration. This setting is primarily used for WAN support.
- **Network adapters**—Select a network adapter from the source and specify the **Replica IP addresses**, **Replica Default Gateways**, and **Replica DNS Server addresses** to be used after failover. If you add multiple gateways or DNS servers, you can sort them by using the arrow up and arrow down buttons. Repeat this step for each network adapter on the source.

Updates made during failover will be based on the network adapter name when protection is established. If you change that name, you will need to delete the job and re-create it so the new name will be used during failover.

If you update one of the advanced settings (IP address, gateway, or DNS server), then you must update all of them. Otherwise, the remaining items will be left blank. If you do not specify any of the advanced settings, the replica virtual machine will be assigned the same network configuration as the source.

By default, the source IP address will be included in the target IP address list as the default address. If you do not want the source IP address to be the default address on the target after failover, remove that address from the **Replica IP addresses** list.

Failover Monitor

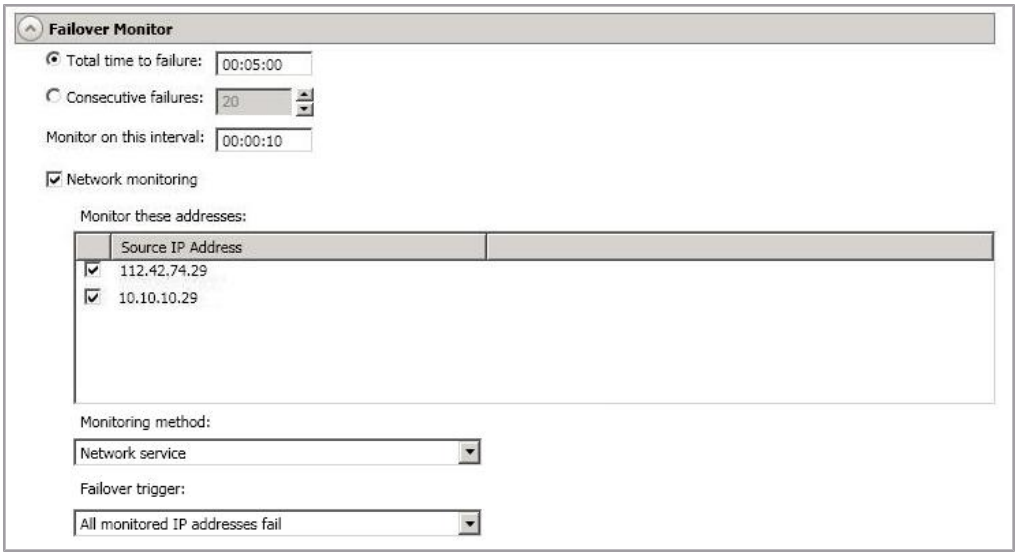

- **Total time to failure**—Specify, in hours:minutes:seconds, how long the target will keep trying to contact the source before the source is considered failed. This time is precise. If the total time has expired without a successful response from the source, this will be considered a failure.

  Consider a shorter amount of time for servers, such as a web server or order processing database, which must remain available and responsive at all times. Shorter times should be used where redundant interfaces and high-speed, reliable network links are available to prevent the false detection of failure. If the hardware does not support reliable communications, shorter times can lead to premature failover. Consider a longer amount of time for machines on slower networks or on a server that is not transaction critical. For example, failover would not be necessary in the case of a server restart.

- **Consecutive failures**—Specify how many attempts the target will make to contact the source before the source is considered failed. For example, if you have this option set to 20, and your source fails to respond to the target 20 times in a row , this will be considered a failure.
- **Monitor on this interval**—Specify, in hours:minutes:seconds, how long to wait between attempts to contact the source to confirm it is online. This means that after a response (success or failure) is received from the source, Double-Take will wait the specified interval time before contacting the source again. If you set the interval to 00:00:00, then a new check will be initiated immediately after the response is received.

  If you choose **Total time to failure**, do not specify a longer interval than failure time or your server will be considered failed during the interval period.

  If you choose **Consecutive failures**, your failure time is calculated by the length of time it takes your source to respond plus the interval time between each response, times the number of consecutive failures that can be allowed. That would be (response time + interval) * failure number. Keep in mind that timeouts from a failed check are included in the response time, so your failure time will not be precise.

- **Network monitoring**—With this option, the target will monitor the source using a network ping.
  - **Monitor these addresses**—Select each **Source IP Address** that you want the target to monitor. If you want to monitor additional addresses, enter the address and click **Add**.
  - **Monitoring method**—This option determines the type of network ping used for failover monitoring.
    - **Network service**—Source availability will be tested by an ICMP ping to confirm the route is active.
    - **Replication service**—Source availability will be tested by a UDP ping to confirm the Double-Take service is active.
    - **Network and replication services**—Source availability will be tested by both an ICMP ping to confirm the route is active and a UDP ping to confirm the Double-Take service is active. Both pings must fail in order to trigger a failover.
  - **Failover trigger**—If you are monitoring multiple IP addresses, specify when you want a failover condition to be triggered.
    - **One monitored IP address fails**—A failover condition will be triggered when any one of the monitored IP addresses fails. If each IP address is on a different subnet, you may want to trigger failover after one fails.
    - **All monitored IP addresses fail**—A failover condition will be triggered when all monitored IP addresses fail. If there are multiple, redundant paths to a server, losing one probably means an isolated network problem and you should wait for all IP addresses to fail.

Failover Options

Wait for user to initiate failover—By default, the failover process will wait for you to initiate it, allowing you to control when failover occurs. When a failure occurs, the job will wait in **Failover Condition Met** for you to manually initiate the failover process. Disable this option only if you want failover to occur immediately when a failure occurs.

Failover Identity

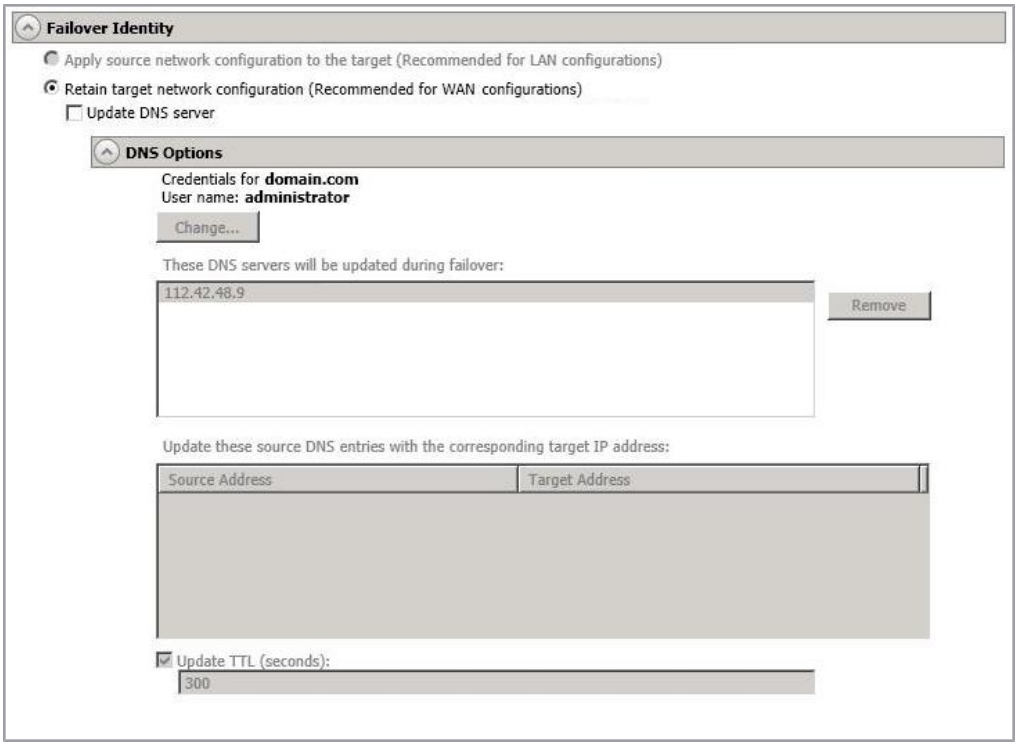

- **Retain target network configuration**—
  - **Update DNS server**—Specify if you want Double-Take to update your DNS server on . If DNS updates are made, the DNS records will be locked during . Be sure and review the *Core Double-Take requirements* on page 23 for the requirements for updating DNS.

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

    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.

If you select **Retain your target network configuration** but do not enable **Update DNS server**, you will need to specify failover scripts that update your DNS server during failover, or you can update the DNS server manually after failover. This would also apply to non--Microsoft Active Directory Integrated DNS servers. You will want to keep your target network configuration but do not update DNS. In this case, you will need to specify failover scripts that update your DNS server during failover, or you can update the DNS server manually after failover.

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

If you are using domain credentials during job creation, you must be able to resolve the domain name from the replica virtual machine using DNS before you can reverse.

Mirror, Verify & Orphaned Files

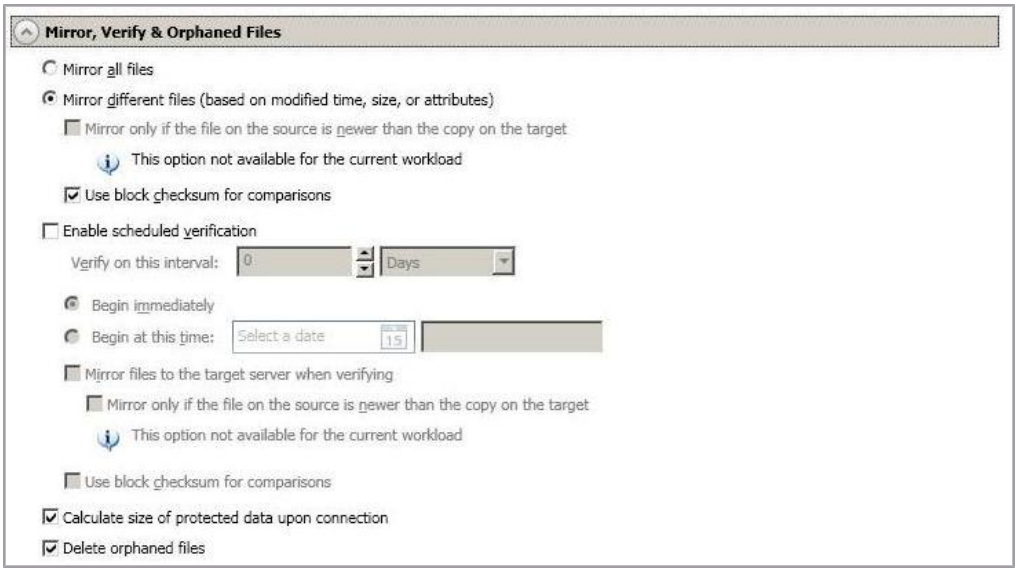

- **Mirror all files**—All protected files will be mirrored from the source to the target.
- **Mirror different files**—Only those protected files that are different based on date and time, size, or attributes will be mirrored from the source to the target.
  - **Mirror only if the file on the source is newer than the copy on the target**—This option is not available for V to Hyper-V jobs.
  - **Use block checksum for comparisons**—For those files flagged as different, the mirroring process can perform a block checksum comparison and send only those blocks that are different.

File Differences Mirror Options Compared

The following table will help you understand how the various difference mirror options work together, including when you are using the block checksum option configured through the *Source server properties* on page 92.

An X in the table indicates that option is enabled. An X enclosed in parentheses (X) indicates that the option can be on or off without impacting the action performed during the mirror.

Not all job types have the source newer option available.

Source Server Properties	Job Properties			Action Performed
Block Checksum Option	File Differences Option	Source Newer Option	Block Checksum Option	
(X)	X			Any file that is different on the source and target based on the date, time, size, and/or attribute is transmitted to the target. The mirror sends the entire file.
(X)	X	X		Any file that is newer on the source than on the target based on date and/or time is transmitted to the target. The mirror sends the entire file.
	X		X	Any file that is different on the source and target based on date, time, size, and/or attributed is flagged as different. The mirror then performs a checksum comparison on the flagged files and only sends those blocks that are different.
X	X		X	The mirror performs a checksum comparison on all files and only sends those blocks that are different.
(X)	X	X	X	Any file that is newer on the source than on the target based on date and/or time is flagged as different. The mirror then performs a checksum comparison on the flagged files and only sends those blocks that are different.

- **Enable scheduled verification**—Verification is the process of confirming that the source replica data on the target is identical to the original 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, can automatically initiate a remirror, if configured. The remirror ensures data integrity between the source and target. When this option is enabled, Double-Take will verify the source replica data on the target and generate a verification log.

  

  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.

  - **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 at the specified date and time.
  - **Mirror files to the target server when verifying**—When this option is enabled, in addition to verifying the data and generating a log, Double-Take will also mirror to the target any protected files that are different on the source.
    - **Mirror only if the file on the source is newer than the copy on the target**—This option is not available for V to Hyper-V jobs.
    - **Use block checksum for comparisons**—For those files flagged as different, the mirroring process can perform a block checksum comparison and send only those blocks that are different.
- **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 the calculation is enabled, it is completed before the job starts mirroring, which can take a significant amount of time depending on the number of files and system performance. If your job contains a large number of files, for example, 250,000 or more, you may want to disable the calculation so that data will start being mirrored sooner. 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 V to Hyper-V jobs. The calculation time is when the system state protection processes hard links. If you disable the calculation, the hard link processing will not occur and you may have problems after failover, especially if your source is Windows 2008.

- **Delete orphaned files**—An orphaned file is a file that exists in the replica data on the

target, but does not exist in the protected data on the source. This option specifies if orphaned files should be deleted on the target during a mirror, verification, or restoration.

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

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 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.

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 95 for more information.

During a mirror, orphaned file processing success messages will be logged to a separate orphaned file log. 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 100 for details on the location of that folder. The orphaned log file is overwritten during each orphaned file processing during a mirror, and the log file will be a maximum of 50 MB.

Network Route

This section is not applicable if your target is a cluster.

For **Send data to the target server using this route**, Double-Take will select, by default, a target route for transmissions. If desired, specify an alternate route on the target that the data will be transmitted through. This allows you to select a different route for Double-Take traffic. For example, you can separate regular network traffic and Double-Take traffic on a machine with multiple IP addresses.

The IP address used on the source will be determined through the Windows route table.

Reverse Network Route

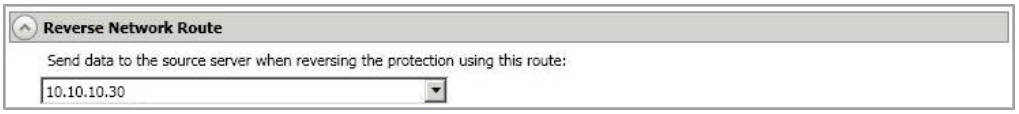

This section is not applicable if your source is a cluster.

Send data to the source server when reversing the protection using this route—By default, Double-Take will select a default source route for transmissions. If desired, specify an alternate route on the source 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.

The IP address used on the target during a reverse will be determined through the Windows route table.

Compression

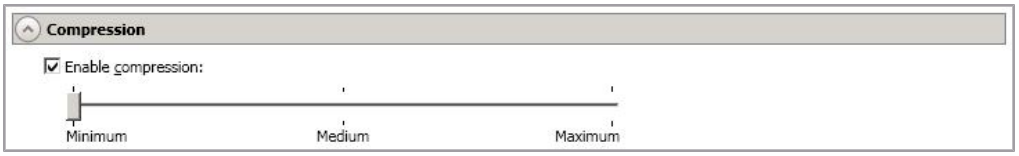

To help reduce the amount of bandwidth needed to transmit Double-Take data, compression allows you to compress data prior to transmitting it across the network. In a WAN environment this provides optimal use of your network resources. If compression is enabled, the data is compressed before it is transmitted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

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

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

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

Bandwidth

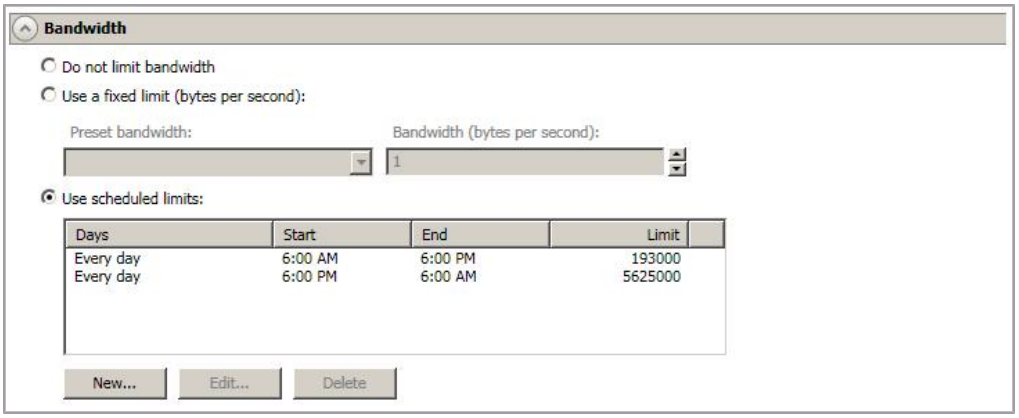

Bandwidth limitations are available to restrict the amount of network bandwidth used for Double-Take data transmissions. When a bandwidth limit is specified, Double-Take never exceeds that allotted amount. The bandwidth not in use by Double-Take is available for all other network traffic.

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

- **Do not limit bandwidth**—Double-Take will transmit data using 100% bandwidth availability.
- **Use a fixed limit**—Double-Take will transmit data using a limited, fixed bandwidth. Select a **Preset bandwidth** limit rate from the common bandwidth limit values. The **Bandwidth** field will automatically update to the bytes per second value for your selected bandwidth. This is the maximum amount of data that will be transmitted per second. If desired, modify the bandwidth using a bytes per second value. The minimum limit should be 3500 bytes per second.
- **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**.

4. If you want to modify the workload items or replication rules for the job, click **Edit workload or replication rules**. Modify the **Workload item** you are protecting, if desired. Additionally, you can modify the specific **Replication Rules** for your job.

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

   You can also enter wildcard rules, however you should do so carefully. Rules are applied to files that are closest in the directory tree to them. If you have rules that include multiple folders, an exclusion rule with a wild card will need to be added for each folder that it needs applied to. For example, if you want to exclude all .log files from D:\ and your rules include D:\, D:\Dir1 , and D:\Dir2, you would need to add the exclusion rule for the root and each subfolder rule. So you will need to add exclude rules for D:*.log , D:\Dir1*.log, and D:\Dir2*.log.

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

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

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

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

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

Viewing a V to Hyper-V job log

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

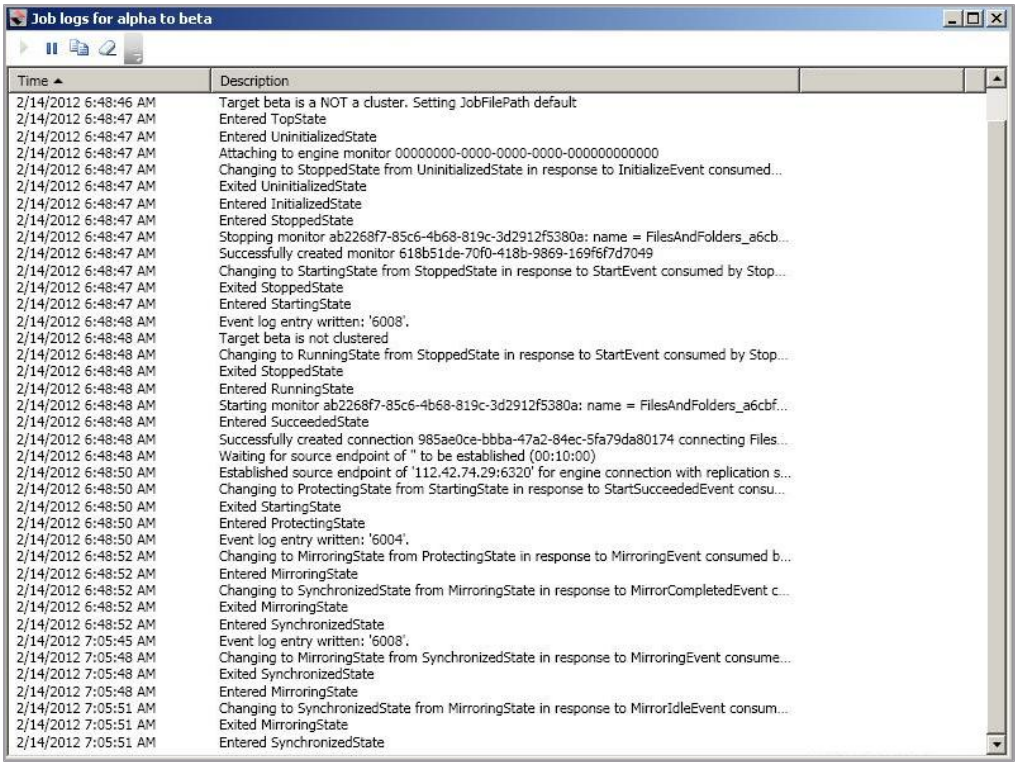

The following table identifies the controls and the table columns in the **Job logs** window.

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.

Failing over V to Hyper-V jobs

When a failover condition has been met, failover will be triggered automatically if you disabled the wait for user option during your failover configuration. If the wait for user before failover option is enabled, you will be notified in the console when a failover condition has been met. At that time, you will need to trigger it manually from the console when you are ready.

1. On the **Manage Jobs** page, highlight the job that you want to failover and click **Failover or Cutover** in the toolbar.
2. Select the type of failover to perform.
   - **Failover to live data**—Select this option to initiate a full, live failover using the current data on the target. This option will shutdown the source machine (if it is online), stop the protection job, and start the replica virtual machine on the target with full network connectivity.
   - **Perform test failover**—Select this option to perform a test failover using the current data on the target. This option will leave the source machine online, stop the protection job, and start the replica virtual machine on the target without network connectivity.
   - **Failover to a snapshot**—This option is not available for V to Hyper-V jobs.
3. Select how you want to handle the data in the target queue. You may want to check the amount of data in queue on the target by reviewing the *Statistics* on page 688 or *Performance Monitor* on page 790.
   - **Apply data in target queues before failover or cutover**—All of the data in the target queue will be applied before failover begins. The advantage to this option is that all of the data that the target has received will be applied before failover begins. 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 failover or cutover immediately**—All of the data in the target queue will be discarded and failover will begin immediately. The advantage to this option is that failover will occur immediately. The disadvantage is that any data in the target queue will be lost.
   - **Revert to last good snapshot if target data state is bad**—If the target data is in a bad state, Double-Take will automatically revert to the last good Double-Take snapshot before failover begins. If the target data is in a good state, Double-Take will not revert the target data. Instead, Double-Take will apply the data in the target queue and then failover. The advantage to this option is that good data on the target is guaranteed to be used. The disadvantage is that if the target data state is bad, you will lose any data between the last good snapshot and the failure.
4. When you are ready to begin failover, click **Failover**.

Depending on your replica configuration, you may have to reboot your replica after failover. You will be prompted to reboot if it is necessary.

In a cluster configuration, if you move the shared storage ownership on the original source cluster or change the drive letter after failover, you will be unable to reverse your protection. Keep the source cluster configuration the same in order to allow proper reversing of protection.

If you used domain credentials when you originally created your job, you must be able to resolve the domain name from the replica virtual machine using DNS before you can reverse.

If your source is running on Windows Server 2008 and the target replica has one or more SCSI drives, then after failover the CD/DVD ROM will not be allocated. If the CD/DVD ROM is required, you will need to edit the virtual machine settings to add a CD/DVD ROM after failover. By not allocating a CD/DVD ROM under these specific conditions, drive letter consistency will be guaranteed.

---

5. If you performed a test failover, you can undo it by selecting **Undo Failover or Cutover** in the toolbar. The replica virtual machine on the target will be shut down and the protection job will be restarted performing a file differences mirror.

Reversing V to Hyper-V jobs

Reversing protection allows you to protect your source replica virtual server running on the target back to the original source host.

1. Make sure you original source virtual machine is not running.
2. If you used domain credentials when you originally created your job, make sure you can resolve the domain name from the replica virtual machine using DNS.
3. On the **Manage Jobs** page, highlight the job that you want to reverse and click **Reverse** in the toolbar. The flow of mirroring and replication data will change. Data will be transmitted from the replica virtual machine on the target back to the original source host.

After the reverse is complete, your source replica virtual machine on the target is being protected to your original source host. In the event you want to go back to your original server roles and hardware configuration, you can failover again.

The original disk structure on the original source virtual machine will be deleted and re-created the first time you perform a reverse. This reverse will perform a full mirror. Subsequent reverses will always perform a file differences mirror.

Chapter 15 Agentless Hyper-V protection

Create an agentless Hyper-V job when you want to protect a virtual machine on a Hyper-V host to another Hyper-V host. This protection type is at the host level. You can reverse the protection after failover has occurred.

- See *Agentless Hyper-V requirements* on page 604—There are specific requirements for agentless Hyper-V protection.
- See *Creating an agentless Hyper-V job* on page 608—This section includes step-by-step instructions for creating an agentless Hyper-V job.
- See *Configuring Hyper-V Pro tip integration for failover notification* on page 631—This section explains how you can extend Microsoft System Center Virtual Machine Manager (SCVMM) Performance and Resource Optimization (PRO) capabilities by providing specific PRO tips for failover of agentless Hyper-V jobs.
- See *Managing and controlling agentless Hyper-V jobs* on page 633—You can view status information about your agentless Hyper-V jobs and learn how to control these jobs.
- See *Failing over agentless Hyper-V jobs* on page 651—Use this section when a failover condition has been met or if you want to failover manually.
- See *Reversing agentless Hyper-V jobs* on page 653—Use this section to reverse protection. The source replica on the target is now sending data back to the original source.

Agentless Hyper-V requirements

In addition to the *Core Double-Take requirements* on page 23, use these requirements for agentless Hyper-V protection.

- **Source and target host operating system**—Your source and target host servers can be any Windows 2008, 2008 R2, 2012, or 2012 R2 operating system from the *Core Double-Take requirements* on page 23 that has the Hyper-V role enabled. In addition, you can use Hyper-V Server 2008 R2, Server Core 2008 R2, Server Core 2012, or Server Core 2012 R2 with the Hyper-V role enabled. (Hyper-V Server 2008 and Server Core 2008 are not supported.) In each case, the source and target must be running identical operating system versions. For example, your source cannot be Windows 2008 (or Windows 2012) and your target Windows 2008 R2 (or Windows 2012 R2).

  

  If your source is using Windows 2008 R2 Service Pack 1 and you the virtual machines you will be protecting are configured to use dynamic memory, then your target must have the same service pack level as the source.

- **Guest operating systems**—The guest operating system can be any operating system. However, if you want Double-Take to monitor the virtual machine for failover, then you must have Integration Components installed on the guest operating system and the virtual machine must be powered on.
- **Server Core**—In addition to the Server Core requirements above, there is a Server Core limitation. DNS updates are not supported for Server Core servers.
- **Virtual machine configurations**—The following limitations apply to the virtual machines on the source and target Hyper-V servers.
  - The virtual machines must be in their own home folder that is not shared by any other virtual machines.
  - The virtual machines cannot be created in or replicated to the Hyper-V system default folder.
  - The virtual machines' snapshot folder must be unique to each virtual machine, they cannot be in the Hyper-V system default folder, and they cannot be changed once protection has been established.
  - The virtual machines cannot use raw, pass-through, or differencing disks.
  - If a source virtual machine is configured for dynamic memory, the replica virtual machine will automatically be configured for dynamic memory. You cannot configure the replica differently.
- **IP addressing**—IPv4 is the only supported IP version.
- **WAN support**—If your source and target are across a WAN and you want Double-Take to automatically update networking on the guest operating system during failover, the following limitations apply. If you choose not to have Double-Take automatically update networking on the guest operating system during failover, you will have to update the network manually, but the following limitations will not apply.

- **Guest operating system**—The guest operating system must be Windows 2003, 2008, or 2012.
- **Windows Management Instrumentation (WMI)**—The host and guest operating systems must have the WMI service enabled.
- **User Access Control (UAC)**—UAC must be disabled on the guest operating system.
- **Name resolution**—You must establish name resolution for the guest operating system.

- **Microsoft .NET Framework**—Microsoft .NET Framework version 3.5 Service Pack 1 is required on the source and target. This version is not included in the .NET version 4.0 release. Therefore, even if you have .NET version 4.0 installed, you will also need version 3.5.1. If you are using Windows 2008 or earlier, you can install this version from the Double-Take DVD, via a web connection during the Double-Take installation, or from a copy you have obtained manually from the Microsoft web site.If you are using Windows 2008 R2 or later, you can enable it through Windows features.
- **Live migration**—Windows 2012 live migration is supported for CSV clustered virtual machines, but the shared-nothing configuration is not supported. If a virtual machine is migrated, Double-Take will automatically start a remirror.
- **Supported configurations**—The following table identifies the supported configurations for an agentless Hyper-V job. (These are source and target configurations for the host, not the virtual machines.)

Configuration		Supported	Not Supported
Source to target configuration[1]	One to one, active/standby	X	
	One to one, active/active	X	
	Many to one	X	
	One to many	X	
	Chained		X
	Single server	X	
Server configuration[2]	Standalone to standalone	X	
	Standalone to cluster	X	
	Cluster to standalone	X	
	Cluster to cluster	X	
	Cluster Shared Volumes (CSV) guest level	X	
	Cluster Shared Volumes (CSV) host level [3,4]	X	

Configuration		Supported	Not Supported
Upgrade configuration[5]	Upgrade 5.3 Hyper-V to Hyper-V job to 7.0 agentless Hyper-V job	X	
	Upgrade 6.0 host-level Hyper-V job to 7.0 agentless Hyper-V job	X	
Version 7.0 console[6,7]	Version 7.0 console can create job for 5.3 source and 5.3 target		X
	Version 7.0 console can create job for 6.0 source and 6.0 target	X	
	Version 7.0 console can create job for 7.0 source and 7.0 target	X	

1. See *Supported configurations* on page 16 for details on each of the source to target configurations.
2. Windows 2008 R2, 2012, and 2012 R2 are the only supported versions of clustered Hyper-V servers. Windows 2008 R1 is not supported for clustering.
3. CSV support at the host-level is for Windows 2012 and 2012 R2 only. Although CSV is not supported for host-level protection for Windows 2008, you can use host-level protection for non-CSV virtual machines on a CSV configured cluster. To do this, you will need to execute the altitude script located in the \tools\scripts\csv directory on the product DVD or in the extracted web download files. If you copy the altitude.bat file to another location, be sure and copy the two .reg files as well. To protect non-CSV virtual machines on a CSV configured cluster, run the following command on each node of the cluster.

   altitude default

   The default script will cause your cluster to go into redirected (maintenance) mode the next time the CSV is brought online, taken offline, or changes owning nodes.

   If you change back to CSV virtual machines on your CSV configured cluster, you will need to run the following command to change Double-Take back to the correct configuration for CSV functionality.

   altitude target

   The target script will keep your cluster from going into redirected (maintenance) mode.

4. If the host containing your source virtual machines is a CSV configuration, the virtual disks must be on the same CSV volume. Different CSV volumes for the source virtual disks is not supported.
5. When upgrading from version 5.3, you can perform a rolling upgrade where you update the target server first. After the upgrade is complete, the source will automatically reconnect to the target. At this point, the job will be an unmanaged job that you can delete or failover. No other job controls will be available. Once you upgrade you source, the job will be fully

controllable. If you are using a cluster configuration, you cannot upgrade. You must delete the existing job, upgrade all of your nodes, and then re-create the job.

6. Once you upgrade your console to version 7.0, existing jobs that are running version 5.3 will not appear in the console until the target of the job is upgraded to version 7.0.
7. Newer job options available in the version 7.0 console will not be functional when creating jobs for servers running version 6.0.

Creating an agentless Hyper-V job

Use these instructions to create an agentless Hyper-V job.

1. Click **Get Started** from the toolbar.
2. Select **Double-Take Availability** and click **Next**.
3. Select **Protect files and folders, an application, or an entire Windows server** and click **Next**.
4. Choose your source server. This is the Hyper-V server or cluster that is hosting the virtual machines that you want to protect.If your virtual machines are on a cluster, select the cluster name, not the node name.

- **Current Servers**—This list contains the servers currently available in your console session. Servers that are not licensed for the workflow you have selected will be filtered out of the list. Select your source server from the list.
- **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 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.

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. The user must also have administrative rights for Microsoft Hyper-V.

5. Click **Next** to continue.
6. Choose the type of workload that you want to protect. Under **Server Workloads**, in the **Workload types** pane, select **Agentless Hyper-V**. In the **Workload items** pane, select the virtual machines on the Hyper-V source that you want to protect. The list of virtual machines will vary depending on whether your source is a Hyper-V server, cluster, or node.

   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.

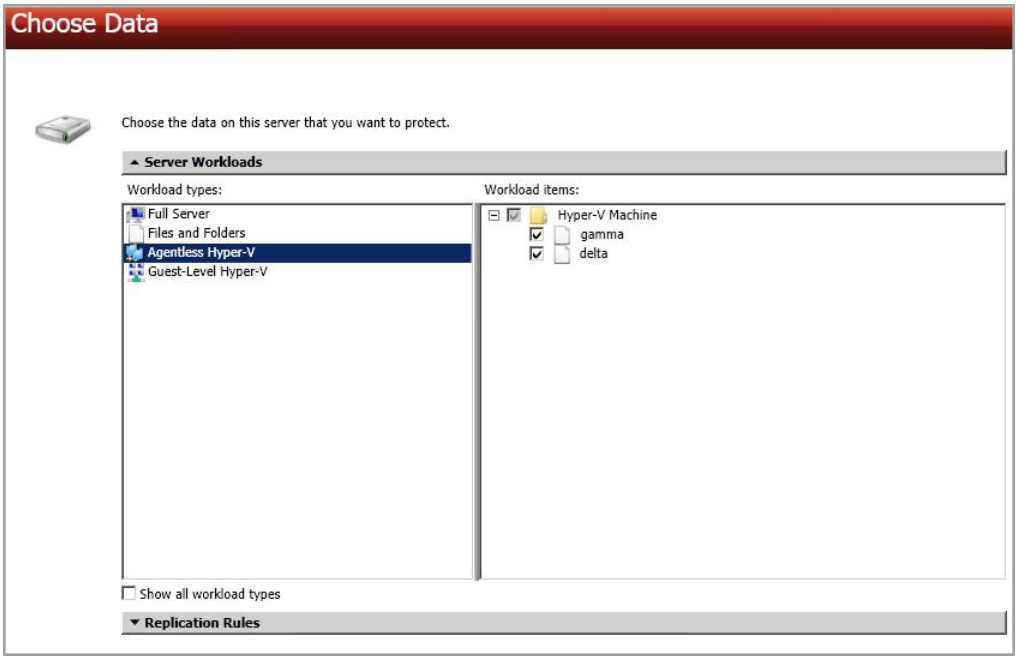

7. Click **Next** to continue.
8. Choose your target server. This is the Hyper-V server or cluster that will store the replicas of the virtual machines from the source.

- **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.
- **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. The user must also have administrative rights for Microsoft Hyper-V.

---

9. Click **Next** to continue.

10. You have many options available for your agentless Hyper-V job. Configure those options that are applicable to your environment.

All agentless Hyper-V jobs will have the following sections available on the **Set Options** page.

- **Replica Virtual Machine Configuration**
- **Mirror, Verify & Orphaned Files**
- **Network Route**
- **Snapshots**
- **Compression**
- **Bandwidth**

If you are protecting just one virtual machine, you will also have the following sections.

- **General**
- **Replica Virtual Machine Location**
- **Replica Virtual Machine Network Settings**
- **Failover Monitor**
- **Failover Identity**

If you are protecting more than one virtual machine, you will have the **Replica Virtual Machines** section instead, which is similar to the location section.

As you can see, if you are protecting more than one virtual machine, there are a few settings that you will not have access to during job creation. In this case, default values will be used. You can modify the default values after the jobs have been created.

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 629.

- *General* on page 612
- *Replica Virtual Machines* on page 613
- *Replica Virtual Machine Location* on page 614
- *Replica Virtual Machine Configuration* on page 615
- *Replica Virtual Machine Network Settings* on page 616
- *Failover Monitor* on page 617
- *Failover Identity* on page 619
- *Mirror, Verify & Orphaned Files* on page 621
- *Network Route* on page 625
- *Snapshots* on page 626
- *Compression* on page 627
- *Bandwidth* on page 628

General

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

Replica Virtual Machines

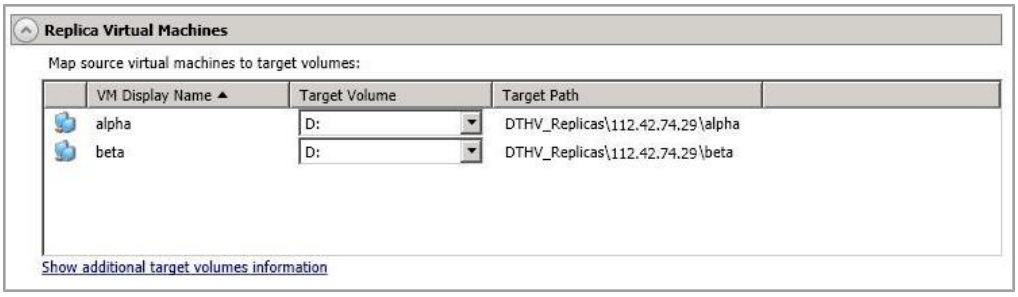

- **Target Volume**—For each virtual machine you are protecting, specify the volume where you want to store the replica virtual machine on the target.
- **Target Path**—For each virtual machine you are protecting, specify a path on the selected **Target Volume** where you want to store the replica virtual machine on the target.
- **Show additional target volumes information**—Click this link to see storage information for the volumes on your target. This will help you select the appropriate volumes for your replica virtual machines.

If your target is a cluster, you will only see the clustered volumes that have been added to the cluster prior to creating the job. If the target is a cluster node, you will only see non-clustered volumes.

Replica Virtual Machine Location

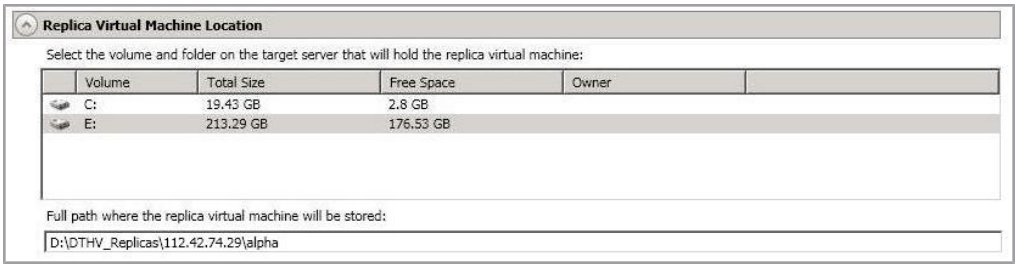

- **Select the volume and folder on the target server that will hold the replica virtual machine**—Select one of the volumes from the list to indicate the volume on the target where you want to store the new virtual server when it is created. The target volume must have enough **Free Space** to store the source data. If your target is a cluster, you will only see the cluster aware volumes that have been added to the cluster prior to creating the job.
- **Full path where the replica virtual machine will be stored**—Specify a location on the selected **Volume** to store the replica of the source. By specifying an existing folder, you can reuse an existing virtual machine on your Hyper-V target created by a previous protection job. This can be useful for pre-staging data on a virtual machine over a LAN connection and then relocating it to a remote site after the initial mirror is complete. You save time by performing a difference mirror instead of a full mirror. In order to use a pre-existing virtual disk, it must be a valid virtual disk and it cannot be attached to any registered virtual machine. In a WAN environment, you may want to take advantage of re-using an existing virtual disk by using a process similar to the following.
    a. Create a protection job in a LAN environment, letting Double-Take create the virtual disk for you.
    b. Complete the mirror process locally.
    c. Delete the protection job and when prompted, select to keep the replica.

    

    Even though you are keeping the replica virtual machine, it is not registered with the Hyper-V Manager. If you want to use the replica machine, you will have to register it. However, you do not need to register it to reuse the disks in Double-Take. Double-Take will automatically register the virtual machine at failover time.

    d. Shut down and move the Hyper-V target server to your remote site.
    e. After the Hyper-V target server is back online at the remote site, create a new protection job for the same source server. Double-Take will reuse the existing hard disk files and perform a difference mirror over the WAN to bring the virtual machine up-to-date.

Replica Virtual Machine Configuration

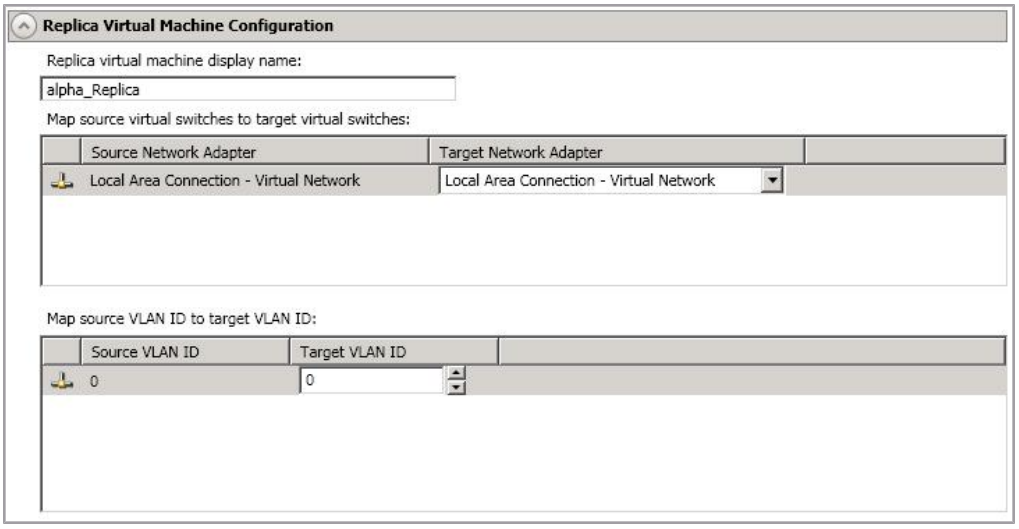

- **Replica virtual machine display name**—Specify the name of the replica virtual machine. This will be the display name of the virtual machine on the host system.
- **Map source virtual switches to target virtual switches**—Identify how you want to handle the network mapping after failover. The **Source Network Adapter** column lists the NICs from the source. Map each one to a **Target Network Adapter**, which is a virtual network on the target. You can also choose to failover the NIC and IP addresses but leave them in a not connected state.
- **Map source VLAN ID to target VLAN ID**—If your environment is a virtual local area network, identify how you want to handle the VLAN IDs after failover. Map each **Source VLAN ID** to a **Target VLAN ID**. If you reconfigure your VLAN IDs after you have created your job, you will need to edit the job to update your VLAN mappings, if you want to use the new configuration.

Replica Virtual Machine Network Settings

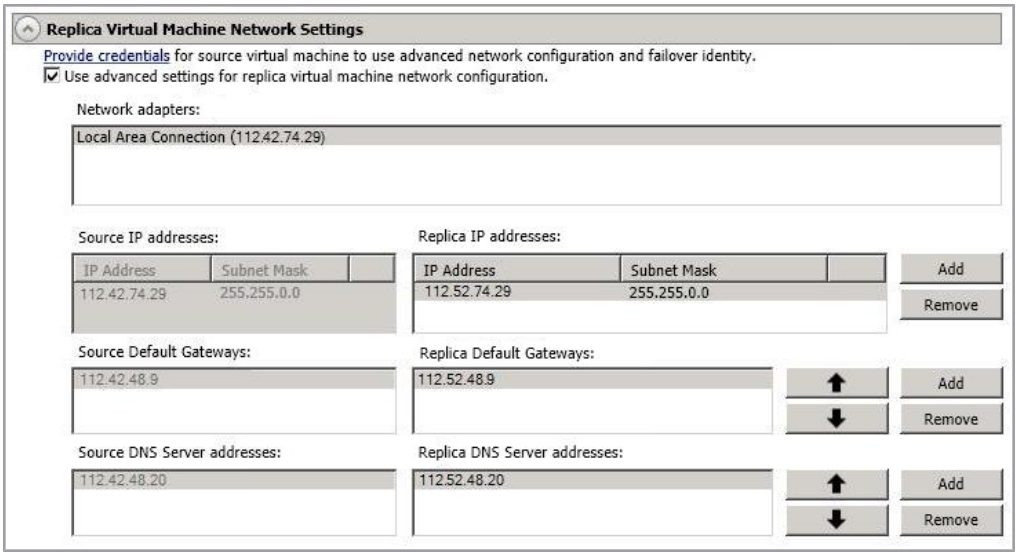

If your virtual machine is powered on and has Integration Services available, this option will allow you to configure advanced settings, which are used primarily for WAN support. Before you can set these options, you must provide credentials for the virtual machine you are protecting. Click the link **Provide credentials** and specify the **Guest Host name**, **User name**, **Password**, **Domain**, and click **OK**

- **Use advanced settings for replica virtual machine network configuration**—Select this option to enable the replica virtual machine network setting configuration.
- **Network adapters**—Select a network adapter from the source and specify the **Replica IP addresses**, **Replica Default Gateways**, and **Replica DNS Server addresses** to be used after failover. If you add multiple gateways or DNS servers, you can sort them by using the arrow up and arrow down buttons. Repeat this step for each network adapter on the source.

Updates made during failover will be based on the network adapter name when protection is established. If you change that name, you will need to delete the job and re-create it so the new name will be used during failover.

If you update one of the advanced settings (IP address, gateway, or DNS server), then you must update all of them. Otherwise, the remaining items will be left blank. If you do not specify any of the advanced settings, the replica virtual machine will be assigned the same network configuration as the source.

By default, the source IP address will be included in the target IP address list as the default address. If you do not want the source IP address to be the default address on the target after failover, remove that address from the **Replica IP addresses** list.

Failover Monitor

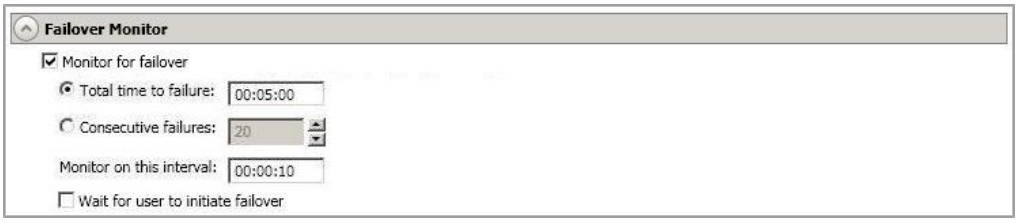

- **Monitor for failover**—Select this option if you want the target to actively monitor the source for a failure. If you disable this option and do not monitor for failover, the target will not actively monitor the source for a failure. In this case, you will have to monitor the source manually on your own and initiate failover manually if there is a source failure.
- **Total time to failure**—Specify, in hours:minutes:seconds, how long the target will keep trying to contact the source before the source is considered failed. This time is precise. If the total time has expired without a successful response from the source, this will be considered a failure.

  Consider a shorter amount of time for servers, such as a web server or order processing database, which must remain available and responsive at all times. Shorter times should be used where redundant interfaces and high-speed, reliable network links are available to prevent the false detection of failure. If the hardware does not support reliable communications, shorter times can lead to premature failover. Consider a longer amount of time for machines on slower networks or on a server that is not transaction critical. For example, failover would not be necessary in the case of a server restart.
- **Consecutive failures**—Specify how many attempts the target will make to contact the source before the source is considered failed. For example, if you have this option set to 20, and your source fails to respond to the target 20 times in a row , this will be considered a failure.
- **Monitor on this interval**—Specify, in hours:minutes:seconds, how long to wait between attempts to contact the source to confirm it is online. This means that after a response (success or failure) is received from the source, Double-Take will wait the specified interval time before contacting the source again. If you set the interval to 00:00:00, then a new check will be initiated immediately after the response is received.

  If you choose **Total time to failure**, do not specify a longer interval than failure time or your server will be considered failed during the interval period.

  If you choose **Consecutive failures**, your failure time is calculated in one of two ways.

  - **ICMP enabled**—If ICMP is enabled on your host when the job is created, your failure time is calculated by the length of time it takes your host to respond to an ICMP host ping plus the length of time it takes your virtual machine to respond plus the interval time between these two checks, times the number of consecutive failures that can be allowed. That would be (ICMP host response time + virtual machine response time + interval) * failure number. If the ICMP check fails, the virtual machine check is skipped and the ICMP check will reoccur after the interval time.
  - **ICMP disabled**—If ICMP is disabled on your host when the job is created, your failure time is calculated by the length of time it takes your virtual machine to respond plus the interval time between each response, times the number of consecutive

failures that can be allowed. That would be (virtual machine response time + interval) * failure number.

Keep in mind that timeouts from a failed check are included in the response times, so your failure time will not be precise.

- **Wait for user to initiate failover**—If you are monitoring for failover, when a failure occurs you can have Double-Take automatically initiate the failover process or wait for you to initiate it. When this option is enabled and a failure occurs, the job will wait in **Failover Condition Met** for you to manually initiate the failover process. When this option is disabled, failover will occur immediately when a failure occurs.

Failover Identity

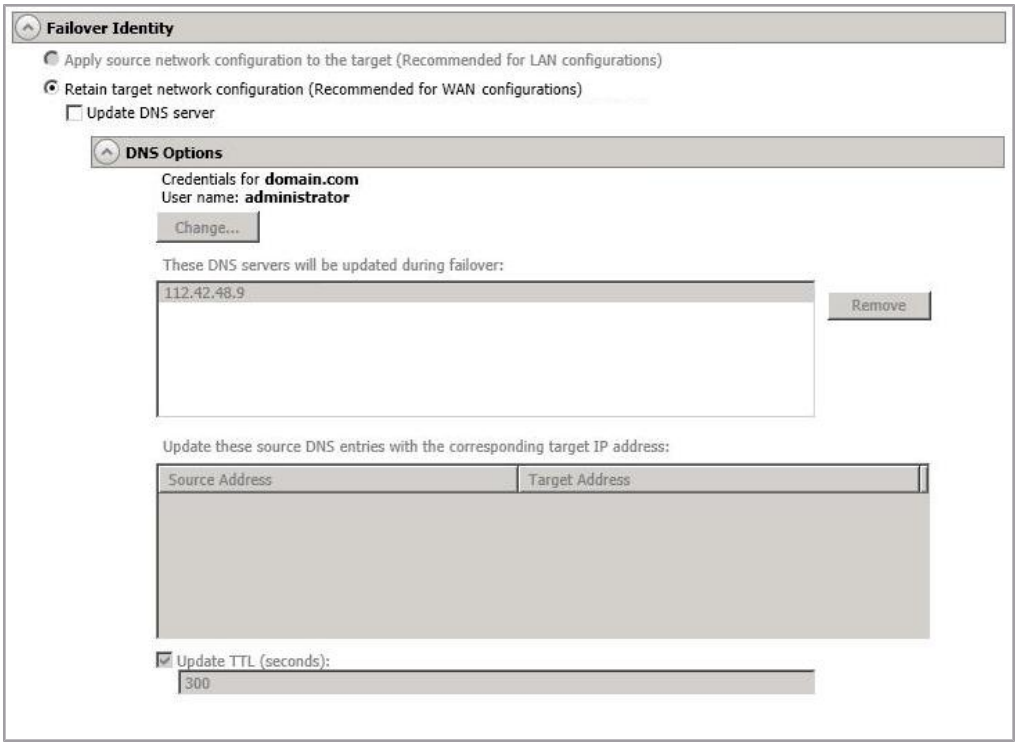

- **Retain target network configuration**—The target will retain all of its original IP addresses.
  - **Update DNS server**—Specify if you want Double-Take to update your DNS server on failover. If DNS updates are made, the DNS records will be locked during failover. Be sure and review the *Core Double-Take requirements* on page 23 for the requirements for updating DNS.

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

    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. For clusters, be sure and select the clustered IP address.
- **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.

If you select **Retain your target network configuration** but do not enable **Update DNS server**, you will need to specify failover scripts that update your DNS server during failover, or you can update the DNS server manually after failover. This would also apply to non--Microsoft Active Directory Integrated DNS servers. You will want to keep your target network configuration but do not update DNS. In this case, you will need to specify failover scripts that update your DNS server during failover, or you can update the DNS server manually after failover.

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

Mirror, Verify & Orphaned Files

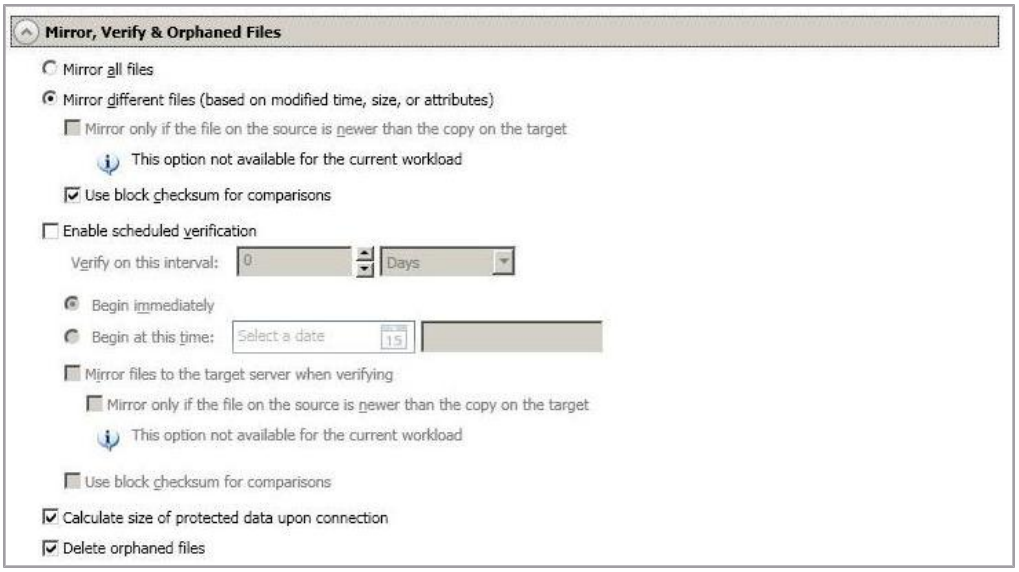

- **Mirror all files**—All protected files will be mirrored from the source to the target.
- **Mirror different files**—Only those protected files that are different based on date and time, size, or attributes will be mirrored from the source to the target.
  - **Mirror only if the file on the source is newer than the copy on the target**—This option is not available for agentless Hyper-V jobs.
  - **Use block checksum for comparisons**—For those files flagged as different, the mirroring process can perform a block checksum comparison and send only those blocks that are different.

File Differences Mirror Options Compared

The following table will help you understand how the various difference mirror options work together, including when you are using the block checksum option configured through the *Source server properties* on page 92.

An X in the table indicates that option is enabled. An X enclosed in parentheses (X) indicates that the option can be on or off without impacting the action performed during the mirror.

Not all job types have the source newer option available.

Source Server Properties	Job Properties			Action Performed
Block Checksum Option	File Differences Option	Source Newer Option	Block Checksum Option	
(X)	X			Any file that is different on the source and target based on the date, time, size, and/or attribute is transmitted to the target. The mirror sends the entire file.
(X)	X	X		Any file that is newer on the source than on the target based on date and/or time is transmitted to the target. The mirror sends the entire file.
	X		X	Any file that is different on the source and target based on date, time, size, and/or attributed is flagged as different. The mirror then performs a checksum comparison on the flagged files and only sends those blocks that are different.
X	X		X	The mirror performs a checksum comparison on all files and only sends those blocks that are different.
(X)	X	X	X	Any file that is newer on the source than on the target based on date and/or time is flagged as different. The mirror then performs a checksum comparison on the flagged files and only sends those blocks that are different.

- **Enable scheduled verification**—Verification is the process of confirming that the source replica data on the target is identical to the original 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, can automatically initiate a remirror, if configured. The remirror ensures data integrity between the source and target. When this option is enabled, Double-Take will verify the source replica data on the target and generate a verification log.

  

  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.

  - **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 at the specified date and time.
  - **Mirror files to the target server when verifying**—When this option is enabled, in addition to verifying the data and generating a log, Double-Take will also mirror to the target any protected files that are different on the source.
    - **Mirror only if the file on the source is newer than the copy on the target**—This option is not available for agentless Hyper-V jobs.
    - **Use block checksum for comparisons**—For those files flagged as different, the mirroring process can perform a block checksum comparison and send only those blocks that are different.
- **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 the calculation is enabled, it is completed before the job starts mirroring, which can take a significant amount of time depending on the number of files and system performance. If your job contains a large number of files, for example, 250,000 or more, you may want to disable the calculation so that data will start being mirrored sooner. 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.

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

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

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 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.

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 95 for more information.

During a mirror, orphaned file processing success messages will be logged to a separate orphaned file log. 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 100 for details on the location of that folder. The orphaned log file is overwritten during each orphaned file processing during a mirror, and the log file will be a maximum of 50 MB.

Network Route

This section is not applicable if your target is a cluster.

For **Send data to the target server using this route**, Double-Take will select, by default, a target route for transmissions. If desired, specify an alternate route on the target that the data will be transmitted through. This allows you to select a different route for Double-Take traffic. For example, you can separate regular network traffic and Double-Take traffic on a machine with multiple IP addresses.

The IP address used on the source will be determined through the Windows route table.

Snapshots

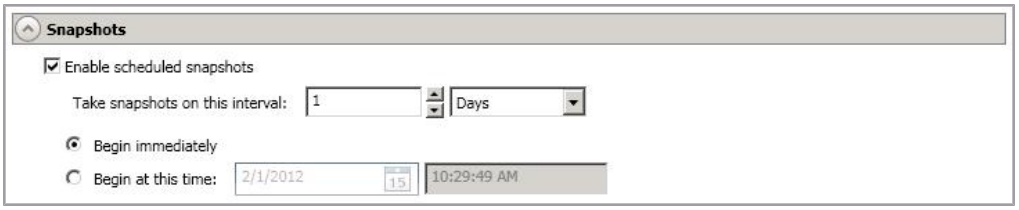

A snapshot is an image of the source replica data on the target taken at a single point in time. You can failover to a snapshot. However, you cannot access the snapshot to recover specific files or folders.

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 starting at a later date and time. Specify the date and time parameters to indicate when you want to start.

See *Managing snapshots* on page 161 for details on taking manual snapshots and deleting snapshots.

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

If your target is a cluster, snapshots will be lost when node ownership changes. However, Double-Take will take an automatic snapshot after a new node becomes owner, so there is always at least one snapshot you can failover to.

Compression

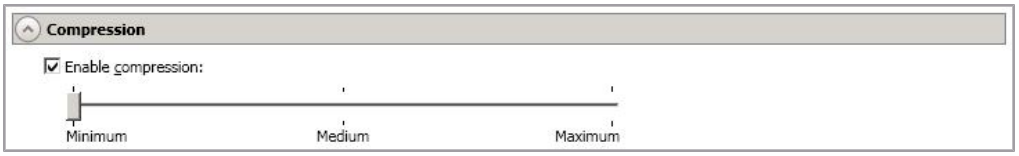

To help reduce the amount of bandwidth needed to transmit Double-Take data, compression allows you to compress data prior to transmitting it across the network. In a WAN environment this provides optimal use of your network resources. If compression is enabled, the data is compressed before it is transmitted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

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

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

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

Bandwidth

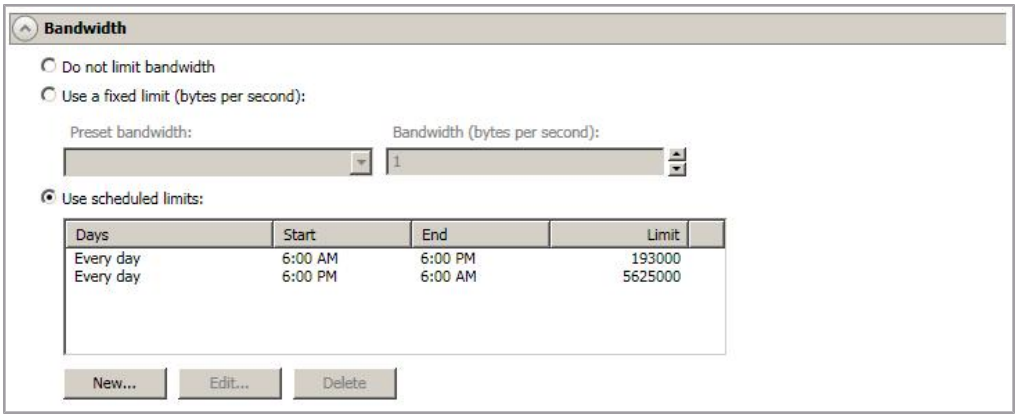

Bandwidth limitations are available to restrict the amount of network bandwidth used for Double-Take data transmissions. When a bandwidth limit is specified, Double-Take never exceeds that allotted amount. The bandwidth not in use by Double-Take is available for all other network traffic.

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

The scheduled option is not available in clustered environments.

- **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**.

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

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

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

Once protection is established, Double-Take monitors the virtual disks of the protected virtual machine for changes to the disk layout. If a new virtual hard disk is added to the virtual machine,

the protection job will automatically be updated to include the new virtual hard disk, and a file difference mirror will automatically start. However, if a virtual hard disk is removed from the protected virtual machine, the virtual hard disk will not be removed from the projection job until it is deleted from the source or the protection job is deleted and re-created.

If your source is a cluster and the Double-Take service on the source is stopped and restarted but failover is not initiated, you will need to manually bring the supporting cluster resource back online through the cluster manager in order for your agentless Hyper-V job to reconnect and transition to a good state. The resource is located on the source cluster and is called DTTargetRes_VM_GUID, where VM is the name of your virtual machine and GUID is a global unique ID assigned to the job. The resource will be located in the Other Resources group.

If your source is a cluster and your cluster resource moves to a node that Double-Take considers a bad node (for example, the Double-Take service is not running or the node has an invalid activation code), your job will enter an error state. You must fix the issue on the cluster node and then you can restart the job. However, if this situation occurs on a target cluster, the job will no longer appear in the console. In this case, you will need to fix the issue on the cluster node and then bring the supporting cluster resource back online through the cluster manager. The resource is located on the source cluster and is called DTTargetRes_VM_GUID, where VM is the name of your virtual machine and GUID is a global unique ID assigned to the job. The resource will be located in the Other Resources group.

Configuring Hyper-V Pro tip integration for failover notification

Microsoft System Center Virtual Machine Manager (SCVMM) provides centralized administration for your Hyper-V virtual machines. Within SCVMM, Performance and Resource Optimization (PRO) provides a basic set of monitors that can alert you to situations where you may want or need to modify a virtual machine configuration in order to optimize the host or virtual machine. These alerts, called PRO tips, recommend actions for you to take to return a host, virtual machine, or any other component of a virtual environment to a healthy state. Double-Take extends the PRO capabilities by providing specific PRO tips for agentless Hyper-V jobs.

In order to receive PRO tips for agentless Hyper-V jobs, your SCVMM machine must be a member of the domain where your Hyper-V host is located, and you must have a machine (same or different than your SCVMM machine) running Microsoft System Center Operations Manager (SCOM) 2007 R2 or 2012.

1. Determine the Double-Take Hyper-V Management Pack for System Center Operations Manager file that you need for your version of SCOM.
   - **SCOM 2007 R2**—Use the file DTHV.FailoverPendingNotification.MP.xml.
   - **SCOM 2012**—Use the file DTHV.FailoverPendingNotification2012.MP.xml.
2. Locate the management pack you need using one of the following methods.
   - Start the Double-Take installation, and when the Autorun appears, select the **Get the SCVMM Management Pack** link. Copy the appropriate .xml file to your SCOM machine.
   - On the Double-Take installation DVD, browse to the SCMgmt\SCVMM directory and copy the appropriate .xml file to your SCOM machine.
   - Download the appropriate Hyper-V Management Pack from the Vision Solutions support web site to your SCOM machine.
3. From the SCOM console, click **Administration** or select **Go**, **Administration**.
4. Right-click on the **Management Packs** line item in the left pane and select **Import Management Pack**.
5. Click **Add** and select **Add from disk**. You can disregard any messages indicating the management pack may have dependencies that cannot be located.
6. Navigate to the location of the Double-Take Hyper-V Management Pack file that you copied or downloaded, and follow the steps in the Management Pack Import Wizard. See the SCOM documentation for complete details.
7. See *Creating an agentless Hyper-V job* on page 608 to configure failover monitoring and manual intervention.
8. In the event your Hyper-V source fails, you will receive a PRO tip alert.

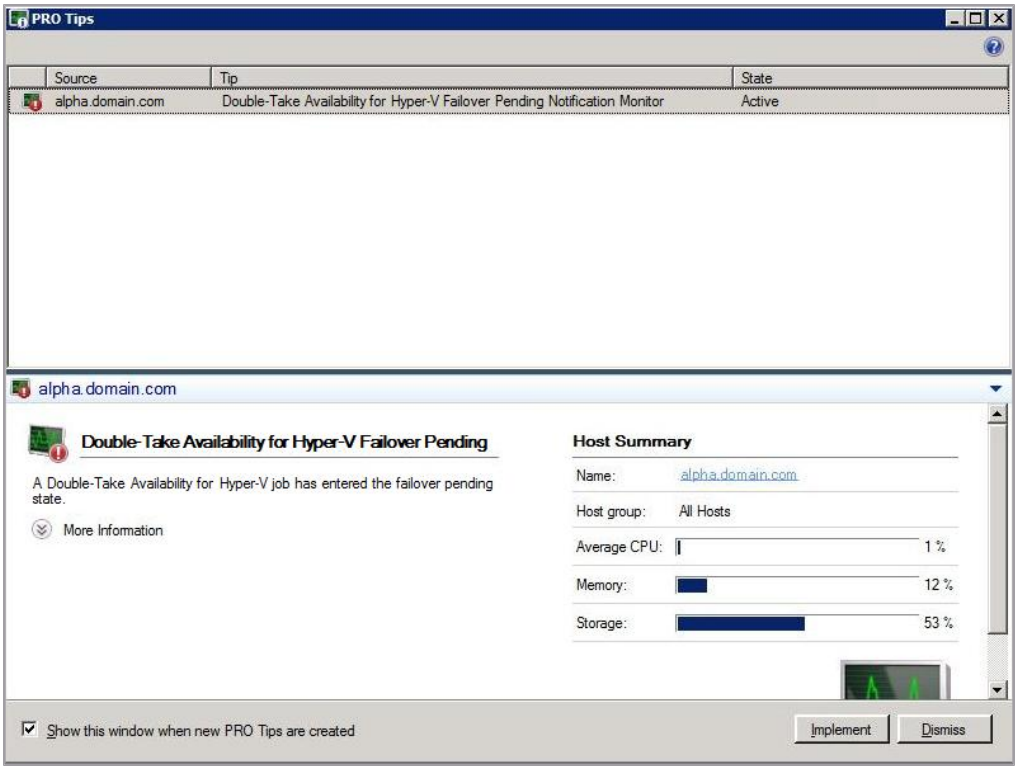

9. You can **Implement** the Pro tip to start failover or you can **Dismiss** it if you do not want to failover.

Managing and controlling agentless Hyper-V jobs

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

The jobs displayed in the right pane depend on the server group folder selected in the left pane. Every job for each server in your console session is displayed when the **Jobs on All Servers** group is selected. If you have created and populated server groups (see *Managing servers* on page 65), then only the jobs associated with the server or target servers in that server group will be displayed in the right pane.

- See *Overview job information displayed in the top pane* on page 633
- See *Detailed job information displayed in the bottom pane* on page 636
- See *Job controls* on page 638

Overview job information displayed in the top pane

The top pane displays high-level overview information about your jobs.

Column 1 (Blank)

The first blank column indicates the state of the job.

 The job is in a healthy state.

 The job is in a warning state. This icon is also displayed on any server groups that you have created that contain a job in a warning state.

 The job is in an error state. This icon is also displayed on any server groups that you have created that contain a job in an error state.

 The job is in an unknown state.

Job

The name of the job

Source Server

The name of the source. This could be a name or IP address of a standalone server, a cluster, or a node. Cluster jobs will be associated with the cluster name and standalone jobs will be associated with a standalone server or a cluster node.

Target Server

The name of the target. This could be a name or IP address of a standalone server, a cluster, or a node. Cluster jobs will be associated with the cluster name and standalone jobs will be associated with a standalone server or a cluster node.

Job Type

Each job type has a unique job type name. This job is an Agentless Hyper-V job. For a complete list of all job type names, press F1 to view the Double-Take Console online help.

Activity

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

If your source is a cluster and the Double-Take service on the source is stopped and restarted but failover is not initiated, you will need to manually bring the supporting cluster resource back online through the cluster manager in order for your agentless Hyper-V job to reconnect and transition to a good state. The resource is located on the source cluster and is called DTTargetRes_VM_GUID, where VM is the name of your virtual machine and GUID is a global unique ID assigned to the job. The resource will be located in the Other Resources group.

If your source is a cluster and your cluster resource moves to a node that Double-Take considers a bad node (for example, the Double-Take service is not running or the node has an invalid activation code), your job will enter an error state. You must fix the issue on the cluster node and then you can restart the job. However, if this situation occurs on a target cluster, the job will no longer appear in the console. In this case, you will need to fix the issue on the cluster node and then bring the supporting cluster resource back online through the cluster manager. The resource is located on the source cluster and is called DTTargetRes_VM_GUID, where VM is the name of your virtual machine and GUID is a global unique ID assigned to the job. The resource will be located in the Other Resources group.

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.
- **Restoring**—Data is being restored from the target to the source.
- **Unknown**—The console cannot determine the status.

Replication Status

- **Replicating**—Data is being replicated to the target.
- **Ready**—There is no data to replicate.

- **Pending**—Replication is pending.
- **Stopped**—Replication has been stopped.
- **Out of Memory**—Replication memory has been exhausted.
- **Failed**—The Double-Take service is not receiving replication operations from the Double-Take driver. Check the Event Viewer for driver related issues.
- **Unknown**—The console cannot determine the status.

Transmit Mode

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

Detailed job information displayed in the bottom pane

The details displayed in the bottom pane of the **Manage Jobs** page provide additional information for the job highlighted in the top pane. If you select multiple jobs, the details for the first selected job will be displayed.

Name

The name of the job

Target data state

- **OK**—The data on the target is in a good state.
- **Mirroring**—The target is in the middle of a mirror process. The data will not be in a good state until the mirror is complete.
- **Mirror Required**—The data on the target is not in a good state because a remirror is required. This may be caused by an incomplete or stopped mirror or an operation may have been dropped on the target.
- **Restore Required**—The data on the source and target do not match because of a failover condition. Restore the data from the target back to the source. If you want to discard the changes on the target, you can remirror to resynchronize the source and target.
- **Snapshot Reverted**—The data on the source and target do not match because a snapshot has been applied on the target. Restore the data from the target back to the source. If you want to discard the changes on the target, you can remirror to resynchronize the source and target.
- **Busy**—The source is low on memory causing a delay in getting the state of the data on the target.
- **Not Loaded**—Double-Take target functionality is not loaded on the target server. This may be caused by an activation code error.
- **Unknown**—The console cannot determine the status.

Mirror remaining

The total number of mirror bytes that are remaining to be sent from the source to the target

Mirror skipped

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

Replication queue

The total number of replication bytes in the source queue

Disk queue

The amount of disk space being used to 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**.

Connected since

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

Recent activity

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

Additional information

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

Job controls

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

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

Create a New Job 

This button leaves the **Manage Jobs** page and opens the **Get Started** page.

View Job Details 

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

Delete 

Stops (if running) and deletes the selected jobs.

If you no longer want to protect the source and no longer need the replica of the source on the target, select to delete the associated replica virtual machine. Selecting this option will remove the job and completely delete the replica virtual machine on the target.

If you no longer want to mirror and replicate data from the source to the target but still want to keep the replica of the source on the target, select to keep the associated replica virtual machine. You may want to use this option to relocate the virtual hard disks and create a new job between the original source and the new location. Selecting this option, will preserve the source replica on the target.

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 77. You will remain on the **Manage Jobs** page after updating the server credentials. If your servers use the same credentials, make sure you also update the credentials on the **Manage Servers** page so that the Double-Take Console can authenticate to the servers in the console session. See *Managing servers* on page 65.

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. Failover monitoring will continue 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. Failover monitoring will continue while the job is stopped.

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 161 for more information.

Failover or Cutover 

Starts the failover process. See *Failing over agentless Hyper-V jobs* on page 651 for the process and details of failing over an agentless Hyper-V job.

Failback 

Starts the failback process. Failback does not apply to agentless Hyper-V jobs.

Restore 

Starts the restoration process. Restore does not apply to agentless Hyper-V jobs.

Reverse 

Reverses protection. The job will start mirroring in the reverse direction with the job name and log file names changing accordingly. After the mirror is complete, the job will continue running in the opposite direction. See *Reversing agentless Hyper-V jobs* on page 653 for the process and details of reversing an agentless Hyper-V job.

Undo Failover 

Cancels failover by undoing it. This resets the servers and the job back to their original state. If you had performed a live failover, any changes made on the target will be lost when you undo. See *Failing over agentless Hyper-V jobs* on page 651 for the process and details of undoing a failed over agentless Hyper-V job.

View Job Log 

Opens the job log. On the right-click menu, this option is called **View Logs**, and you have the option of opening the job log, source server log, or target server log. See *Viewing the log files through the Double-Take Console* on page 678 for details on all three of these logs.

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 all files**—All protected files will be mirrored from the source to the target.

- **Mirror different files**—Only those protected files that are different based on date and time, size, or attributes will be mirrored from the source to the target.
  - **Mirror only if the file on the source is newer than the copy on the target**—This option is available for agentless Hyper-V jobs, but ideally it should not be used.
  - **Use block checksum for comparisons**—For those files flagged as different, the mirroring process can perform a block checksum comparison and send only those blocks that are different.
- **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 the calculation is enabled, it is completed before the job starts mirroring, which can take a significant amount of time depending on the number of files and system performance. If your job contains a large number of files, for example, 250,000 or more, you may want to disable the calculation so that data will start being mirrored sooner. 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.
  - **Create verification report only**—This option verifies the data and generates a verification log, but it does not remirror any files that are different on the source and target. See *Verification log* on page 102 for details on the log file.
  - **Mirror files to the target server automatically**—When this option is enabled, in addition to verifying the data and generating a log, Double-Take will also mirror to the target any protected files that are different on the source.
    - **Mirror only if the file on the source is newer than the copy on the target**—This option is available for agentless Hyper-V jobs, but ideally it should not be used.
    - **Use block checksum for comparisons**—For those files flagged as different, the mirroring process can perform a block checksum comparison and send only those blocks that are different.
- **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.
- **Update Shares**—Shares are not applicable because they are automatically included with the system state that is being protected with the entire server.

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 65.

Type a server name

Displays only jobs that contain the text you entered. If you have created and populated server groups, then only jobs that contain the text you entered associated with the server or target servers in that server group will be displayed. See *Managing servers* on page 65.

Overflow Chevron 

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

Viewing agentless Hyper-V job details

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

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

Job name

The name of the job

Job type

Each job type has a unique job type name. This job is an Agentless Hyper-V job. For a complete list of all job type names, press F1 to view the Double-Take Console online help.

Health

 The job is in a healthy state.

 The job is in a warning state.

 The job is in an error state.

 The job is in an unknown state.

Activity

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

Connection ID

The incremental counter used to number connections. The number is incremented when a connection is created. It is also incremented by internal actions, such as an auto-disconnect and auto-reconnect. The lowest available number (as connections are created, stopped, deleted, and so on) will always be used. The counter is reset to one each time the Double-Take service is restarted.

Transmit mode

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

Target data state

- **OK**—The data on the target is in a good state.
- **Mirroring**—The target is in the middle of a mirror process. The data will not be in a good state until the mirror is complete.
- **Mirror Required**—The data on the target is not in a good state because a remirror is required. This may be caused by an incomplete or stopped mirror or an operation may have been dropped on the target.
- **Restore Required**—The data on the source and target do not match because of a failover condition. Restore the data from the target back to the source. If you want to discard the changes on the target, you can remirror to resynchronize the source and target.
- **Snapshot Reverted**—The data on the source and target do not match because a snapshot has been applied on the target. Restore the data from the target back to the source. If you want to discard the changes on the target, you can remirror to resynchronize the source and target.
- **Busy**—The source is low on memory causing a delay in getting the state of the data on the target.
- **Not Loaded**—Double-Take target functionality is not loaded on the target server. This may be caused by an activation code error.
- **Unknown**—The console cannot determine the status.

Target route

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

Compression

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

Encryption

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

Bandwidth limit

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

Connected since

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

Additional information

Depending on the current state of your job, you may see additional information displayed to keep you informed about the progress and status of your job. If there is no

additional information, you will see (None) displayed.

Mirror status

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

Mirror percent complete

The percentage of the mirror that has been completed

Mirror remaining

The total number of mirror bytes that are remaining to be sent from the source to the target

Mirror skipped

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

Replication status

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

Replication queue

The total number of replication bytes in the source queue

Disk queue

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

Bytes sent

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

Bytes sent compressed

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

Validating an agentless Hyper-V job

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

1. From the **Manage Jobs** page, highlight the job and click **View Job Details** in the toolbar.
2. In the **Tasks** area on the right on the **View Job Details** page, click **Validate job properties**.
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 enable protection. Depending on the error, you may be able to click **Fix** or **Fix All** and let Double-Take correct the problem for you. For those errors that Double-Take cannot correct automatically, you will need to modify the source or target to correct the error, or you can select a different target. You must revalidate the selected servers, by clicking **Recheck**, until the validation check passes without errors.

   Validation checks for an existing job are logged to the job log on the target server. See *Log files* on page 677 for details on the various log files.
4. Once your servers have passed validation, click **Close**.

Editing an agentless Hyper-V job

Use these instructions to edit an agentless Hyper-V job.

1. From the **Manage Jobs** page, highlight the job and click **View Job Details** in the toolbar.
2. In the **Tasks** area on the right on the **View Job Details** page, click **Edit job properties**. (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 agentless Hyper-V 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 an agentless Hyper-V job* on page 608 for details on each job option.

   

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

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

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

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

Viewing an agentless Hyper-V job log

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

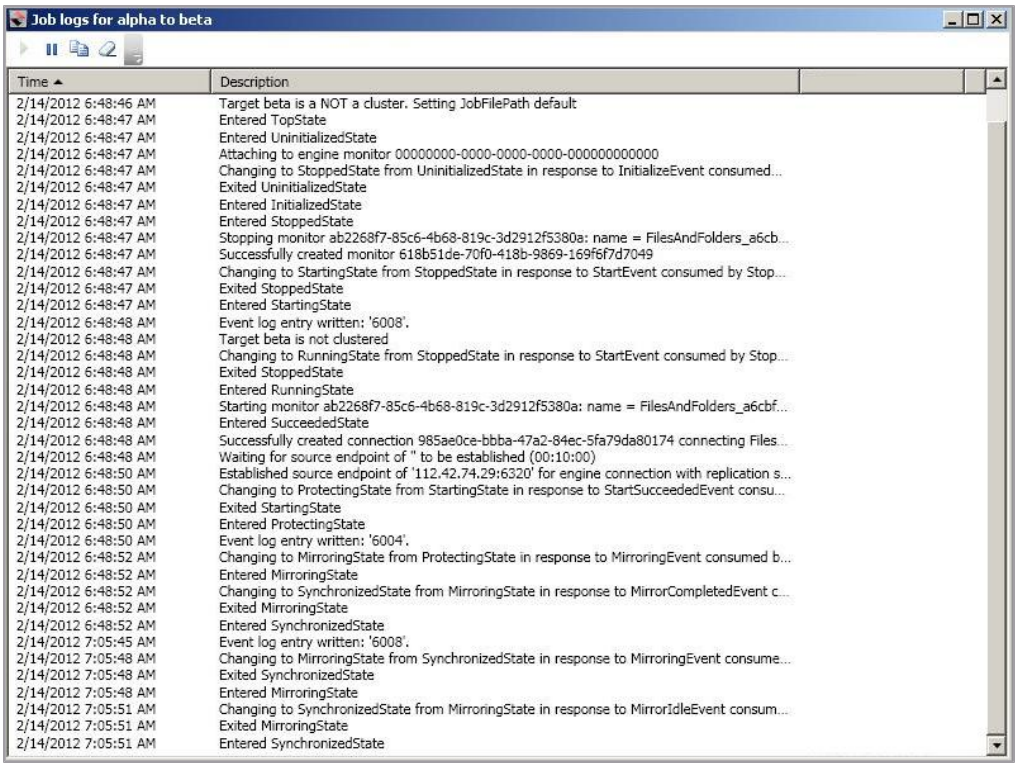

The following table identifies the controls and the table columns in the **Job logs** window.

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.

Failing over agentless Hyper-V jobs

When a failover condition has been met, failover will be triggered automatically if you disabled the wait for user option during your failover configuration. If the wait for user before failover option is enabled, you will be notified in the console when a failover condition has been met. At that time, you will need to trigger it manually from the console when you are ready.

1. On the **Manage Jobs** page, highlight the job that you want to failover and click **Failover or Cutover** in the toolbar.
2. Select the type of failover to perform.
    - **Failover to live data**—Select this option to initiate a full, live failover using the current data on the target. This option will shutdown the source machine (if it is online), stop the protection job, and start the replica virtual machine on the target with full network connectivity.
    - **Perform test failover**—Select this option to perform a test failover using the current data on the target. This option will leave the source machine online, stop the protection job, and start the replica virtual machine on the target without network connectivity.
    - **Failover to a snapshot**—Select this option to initiate a full, live failover without using the current data on the target. Instead, select a snapshot and the data on the target will be reverted to that snapshot. This option will not be available if there are no snapshots on the target or if the target does not support snapshots. This option is also not applicable to clustered environments. To help you understand what snapshots are available, the **Type** indicates the kind of snapshot.
        - **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. Select how you want to handle the data in the target queue. You may want to check the amount of data in queue on the target by reviewing the *Statistics* on page 688 or *Performance Monitor* on page 790.
    - **Apply data in target queues before failover or cutover**—All of the data in the target queue will be applied before failover begins. The advantage to this option is that all of the data that the target has received will be applied before failover begins. 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 failover or cutover immediately**—All of the data in the target queue will be discarded and failover will begin immediately. The advantage to this option is that failover will occur immediately. The disadvantage is that any data in the target queue will be lost.
    - **Revert to last good snapshot if target data state is bad**—If the target data is in a bad state, Double-Take will automatically revert to the last good Double-Take snapshot before failover begins. If the target data is in a good state, Double-Take will not revert the target data. Instead, Double-Take will apply the data in the target queue and then failover. The advantage to this option is that good data on the target is guaranteed to be used. The

disadvantage is that if the target data state is bad, you will lose any data between the last good snapshot and the failure.

4. When you are ready to begin failover, click **Failover**.

Depending on your replica configuration, you may have to reboot your replica after failover. You will be prompted to reboot if it is necessary.

In a cluster configuration, if you move the shared storage ownership on the original source cluster or change the drive letter after failover, you will be unable to reverse your protection. Keep the source cluster configuration the same in order to allow proper reversing of protection.

5. If desired, you can undo your live or test failover by selecting **Undo Failover or Cutover** in the toolbar. In either case, the replica virtual machine on the target will be shut down and the protection job will be restarted performing a file differences mirror. The one difference when undoing a live failover is the virtual machine on the source will be started. This step is not needed when undoing a test failover, because the virtual machine on the source is not shut down during a test failover. In both cases, all changes made on the replica virtual machine on the target will be lost. If you do not want to lose data changes made on the replica virtual machine on the target, see *Reversing agentless Hyper-V jobs* on page 653

Reversing agentless Hyper-V jobs

Reversing protection allows you to protect your source replica virtual server running on the target back to the original source host.

1. Make sure you original source virtual machine is not running.
2. On the **Manage Jobs** page, highlight the job that you want to reverse and click **Reverse** in the toolbar. The flow of mirroring and replication data will change. Data will be transmitted from the replica virtual machine on the target back to the original source host.

After the reverse is complete, your source replica virtual machine on the target is being protected to your original source host. In the event you want to go back to your original server roles and hardware configuration, you can failover again.

Chapter 16 GeoCluster protection

Create a GeoCluster job when you want to use a Microsoft Cluster but do not have shared storage. Review the *GeoCluster requirements* on page 655 and then proceed with your GeoCluster protection using the following steps, in order.

1. Configure a cluster that does not require shared storage. This configuration includes:
   - Creating the cluster
   - Adding nodes to the cluster
   - Setting the cluster quorum
   - Installing Double-Take

   See *Configuring a cluster for GeoCluster* on page 656 for the correct order of these cluster configuration steps, which is dependent on your operating system.
2. Create a new cluster group using the option **Create an empty service or application**. This will be the group name for the application.
3. Create a GeoCluster Replicated Disk resource in your application group. See *Creating a GeoCluster Replicated Disk resource* on page 659.
4. Install your application specifying the application group for the server, database instance, or name that your application requires. For example, if you were using Microsoft SQL Server 2008 R2, you would use the following installation procedure.
   a. SQL 2008 installation on the first node
      1. Select **New SQL Server failover cluster installation**.
      2. Select **Features to install**.
      3. Provide the SQL Server Network Name and Instance ID.
      4. Select the disk resource, which is the GeoCluster Replicated Disk resource that you created.
      5. Provide an IP address for the network resource.
      6. Complete the remaining installation steps using your SQL Server documentation.
   b. SQL 2008 installation on additional node(s)
      a. Select Add node to SQL Server failover cluster.
      b. For the cluster configuration, select the SQL Server instance name from the installation on the first node.
      c. Complete the remaining installation steps using your SQL Server documentation.
5. After the application installation is complete, edit the properties of the resources for your application and make them dependent on the GeoCluster Replicated Disk resource, if needed. This will ensure that the replicated data (as opposed to shared storage) is available before your application starts. In the SQL example, the installation will automatically set the dependencies.

GeoCluster requirements

Make sure your cluster meets the *Core Double-Take requirements* on page 23, as well as the following requirements specific to GeoCluster protection.

- **Operating system**—The operating systems are limited to the Windows Standard (2012 only), Enterprise, and Datacenter editions listed in the *Core Double-Take requirements* on page 23.
- **Hardware**—Microsoft support for MSCS and MSCS-based Microsoft applications requires that the cluster configuration appear on the Microsoft Hardware Compatibility List under category Cluster.
- **Cluster Network Name**—Double-Take does not handle dynamic changes to the cluster network names, the names assigned to the routes for network traffic. If a network name is changed for a network that is used by Double-Take, the GeoCluster Replicated Disk resource must be taken offline, the resource's network property must be changed, and then the resource must be brought back online.
- **Disk queuing**—The Double-Take disk queue, configured during installation, should use a local volume for each node in the cluster.
- **Anti-virus software**—You should configure your anti-virus software to delete or quarantine viruses because cleaning them can cause an access denied retrying operation error. Additionally, configuring virus software to scan outgoing traffic will lessen performance impacts.
- **Supported upgrade configurations**—Before upgrading be sure and read the *Installation notes* on page 35 and review the proper procedure for upgrading Double-Take on a cluster.

Upgrade Configuration	Supported	Not Supported
Upgrade 5.3 GeoCluster Replicated Disk Resource to 7.0 GeoCluster Replicated Disk Resource	X	
Upgrade 6.0 GeoCluster Replicated Disk Resource to 7.0 GeoCluster Replicated Disk Resource	X	

Configuring a cluster for GeoCluster

Complete the Geocluster installation and configuration appropriate for the operating system you are using.

- *Configuring a Windows 2003 cluster* on page 656
- *Configuring a Windows 2008 or 2012 cluster* on page 657

Configuring a Windows 2003 cluster

In a typical Windows 2003 MSCS shared disk cluster configuration, the quorum resource, by default, is the Local Quorum and is located on the first shared disk in the cluster. Because in a GeoCluster configuration there is no shared physical disk, the Local Quorum will not work as the quorum resource. You will need to choose one of the other Windows quorums. The recommended quorum resource for GeoCluster is the Majority Node Set or Majority Node Set with File Share Witness.

If you are upgrading from a previous GeoCluster version and were using GeoCluster as a quorum, you must select another quorum type. GeoCluster can no longer be used as a quorum resource.

- **Local Quorum**—This quorum is for single node clusters and shared disk clusters. It cannot be used in a GeoCluster configuration.
- **Majority Node Set**—This quorum is for clusters with three or more nodes.
- **Majority Node Set with File Share Witness**—This quorum is for clusters with only two nodes. If you are using Windows 2003 Service Pack 1, see the Microsoft support article 921181 for an update for the File Share Witness. If you are using Service Pack 2 or later, the update is not needed.

Use the following instructions as a guideline for configuring your Windows 2003 cluster. See your Windows cluster documentation as a complete reference.

1. Login with an account that has administrative rights on the domain and the local machine.
2. Create the cluster on the first node, if it is not already created. See your Windows documentation for instructions on how to create a cluster.
3. Add your additional nodes to the cluster. See your Windows documentation for instructions on how to add nodes to the cluster.
4. Install Double-Take Availability on each node of the cluster. See *Installing using the installation wizard* on page 37.
5. Configure your quorum. See your Windows documentation for instructions on configuring the quorum appropriate for your environment.
6. If desired, you can install Double-Take Availability on non-clustered client machines if you want to use Cluster Administrator to control the GeoCluster resources. Install Double-Take Availability, selecting the **Client Components Only** installation option.

Configuring a Windows 2008 or 2012 cluster

The default quorum resource in a Windows 2008 or 2012 environment will vary depending on your configuration (number of nodes, shared disks, and so on). The recommended quorum resource for GeoCluster is the Node and File Share Majority. There are other quorum types available. Review the following list to determine which quorum is appropriate for your environment.

- **Node Majority**—This quorum is recommended for clusters with an odd number of nodes. The cluster can handle failures of half of the nodes (rounding up) minus one and still stay online.
- **Node and Disk Majority**—This quorum is recommended for clusters with an even number of nodes. The cluster can handle failures of half of the nodes (rounding up), as long as the witness disk remains online, and still stay online. If the witness disk fails, the cluster can handle failures of only half of the nodes (rounding up) minus one and still stay online.
- **Node and File Share Majority**—This quorum is recommended for clusters with special configurations, such as GeoCluster. The cluster can handle failures of half of the nodes (rounding up), as long as the witness share remains online, and still stay online. If the witness share fails, the cluster can handle failures of only half of the nodes (rounding up) minus one and still stay online.
- **No Majority: Disk Only**—This quorum is not usually recommended. The cluster can handle failures of all nodes except one and still stay online.

Use the following instructions as a guideline for configuring your Windows 2008 or 2012 cluster. See your Windows cluster documentation as a complete reference.

1. Login with an account that has administrative rights on the domain and the local machine.
2. Create the cluster, if it is not already created. See your Windows documentation for instructions on how to create a cluster.
3. Configure a Node and File Share Majority quorum. See your Windows documentation for instructions on how to configure the quorum.
4. If you are going to be using Hyper-V, install the Hyper-V server role on all nodes in the cluster. Make sure that you have the required Microsoft hotfixes applied, including KB958065 which is a failover clustering hotfix and KB950050.
5. Install Double-Take Availability on each node of the cluster. See *Installing using the installation wizard* on page 37.
6. If desired, you can install Double-Take Availability on non-clustered client machines if you want to use Cluster Administrator to control the GeoCluster resources. Install Double-Take Availability, selecting the **Client Components Only** installation option.
7. If you are going to be using Hyper-V, create your virtual machine from within Hyper-V. Be sure to leave the virtual machine off.
8. From Failover Cluster Management, create your application group or role.

If you are creating a file server using clustered file shares, the path for the file share in the Failover Cluster Management wizard is case-sensitive. If the drive letter is uppercase, the path in the clustered file share wizard must also be uppercase. If the case does not match, the wizard will fail stating the path does not exist.

If your application requires a disk before installation can begin, create an Empty Service or Application or Empty Role. After your GeoCluster Replicated Disk resource is created,

you can delete the empty item and Double-Take will automatically move the GeoCluster Replicated Disk resource to available storage for your application installation.

9. If you are using Hyper-V, add your virtual machine resource to the group or role. Any warnings about storage may be disregarded because the GeoCluster Replicated Disk will alleviate storage requirements.

Creating a GeoCluster Replicated Disk resource

The GeoCluster Replicated Disk resource allows for the real-time copy of data to be available on other nodes in the cluster. In the event of a failure and another node takes ownership, the GeoCluster Replicated Disk resource is also moved to the other node and it continues to replicate data, in real-time, to the remaining nodes in the cluster.

The instructions for creating this resource are different depending on your operating system.

- See *Creating the GeoCluster Replicated Disk Resource on Windows 2003* on page 659
- See *Creating the GeoCluster Replicated Disk Resource on Windows 2008 or 2012* on page 661
- See *Creating the GeoCluster Replicated Disk Resource on Windows 2008 or 2012 Hyper-V* on page 662
- See *Bringing the resource online* on page 664
- *Taking the resource offline* on page 664

Creating the GeoCluster Replicated Disk Resource on Windows 2003

1. From the Failover Cluster Manager, right-click the group that you want to add a replicated disk to and select **New**, **Resource**.
2. Specify the following fields on the New Resource dialog box.
    - **Name**—Specify a name that identifies which application, file set, disk, and so on that you are protecting. This name must be unique within the cluster.
    - **Description**—You can optionally add a more detailed description for this resource.
    - **Resource type**—Specify GeoCluster Replicated Disk.
    - **Group**—The group that you originated the new resource from will be selected. Verify that this is the correct group. If it is not, select the correct group name.
3. Click **Next** to continue.
4. The GeoCluster Replicated Disk resource ensures that an up-to-the-minute copy of the data resides on all nodes identified in the **Possible owners** list. All nodes are included in the default, which should not be changed. Click **Next** to continue.
5. The GeoCluster Replicated Disk resource is not dependent on any other resources. Click **Next** to continue.
6. Specify the GeoCluster Replicated Disk parameters using the settings below.
    - **Disk to replicate**—Select a disk to replicate from the available volumes. The only volumes that will be displayed are those that meet the following criteria.
        - NTFS volumes
        - Volumes which are not already being replicated by another GeoCluster Replicated Disk resource
        - Volumes that are not physical disk resources
        - Volumes that do not contain system files (The volume that you booted Windows from will not be displayed.)
        - Volumes that exist on all nodes of the cluster

- **Network**—Select the network that you want to use for Double-Take mirroring and replication traffic. If you do not have multiple networks established, you will only be able to select the one network that does exist. If you do not select a network, Double-Take will use DNS to determine a network route to use. Ideally, the networks used for various traffic should be separated. This is dependent on the number of networks that you established when you created the cluster and the priority assigned to each network. For example, if you have two network routes, separate Double-Take and your public traffic. If you have three routes, separate the public traffic and then separate Double-Take from the cluster heartbeat.
- **Interval to check unresponsive nodes**—Specify how much time, in seconds, between checks of nodes to see if a Double-Take connection can be made.
- **Delay connection until resources dependent on this one are online**—This option allows you to delay a Double-Take connection until any resources that have the GeoCluster Replicated Disk resource as a dependency are online. By ensuring that all resources that are dependent on the GeoCluster Replicated disk resource are online before starting the connection, the chance of a conflict occurring because application resources are attempting to open files exclusively while Double-Take is mirroring those files is removed.

7. Click **Next** to continue.
8. An orphan is a file that exists in the target location but is not in the source location. You can select to delete orphans during a mirror.
9. Click **Next** to continue.
10. If you want to configure compression, verify that **Enable Compression** is selected. Depending on the compression algorithms available for your operating system, you may see a slider bar indicating different compression levels. Set the level from minimum to maximum compression to suit your needs.
11. Click **Next** to continue.
12. Specify your Double-Take mirroring settings.
    - **Full Mirror**—All files will be sent from the source to the target.
    - **File differences**—Only those files that are different based size or date and time will be sent from the source to the target.
      - **Send data only if Source is newer than Target**—Only those files that are newer on the source are sent to the target.

        

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

      - **Use block checksum**—For those files flagged as different, the mirror performs a checksum comparison and only sends those blocks that are different.
    - **Calculate Replication Set size prior to mirror**—Determines the size of the protected data set prior to starting the mirror. The mirroring status will update the percentage complete if the data set size is calculated.
13. Click **Finish** to complete the creation of the GeoCluster Replicated Disk resource.
14. To control the resource, you can bring it online and take it offline. Neither of these actions trigger failover. They just control the activity of the resource.

Creating the GeoCluster Replicated Disk Resource on Windows 2008 or 2012

1. From the Failover Cluster Manager, right-click the application group or role where you want to add a replicated disk to and select **Add a resource**, **More resources**, **Add GeoCluster Replicated Disk** (for Windows 2008) or **GeoCluster Replicated Disk** (for Windows 2012).
2. Right-click on the new resource and select **Properties**.
3. On the **General** tab, specify a name that identifies which application, file set, disk, and so on that you are protecting. This name must be unique within the cluster.
4. On the **Connection Parameters** tab, specify the GeoCluster Replicated Disk connection parameters using the settings below.
    - **Disk to replicate**—Select a disk to replicate from the available volumes. The only volumes that will be displayed are those that meet the following criteria.
        - NTFS volumes
        - Volumes which are not already being replicated by another GeoCluster Replicated Disk resource
        - Volumes that are not physical disk resources
        - Volumes that do not contain system files (The volume that you booted Windows from will not be displayed.)
        - Volumes that exist on all nodes of the cluster
    - **Network to route Double-Take mirroring and replication traffic over**—Select the network that you want to use for Double-Take mirroring and replication traffic. If you do not have multiple networks established, you will only be able to select the one network that does exist. If you do not select a network, Double-Take will use DNS to determine a network route to use. Ideally, the networks used for various traffic should be separated. This is dependent on the number of networks that you established when you created the cluster and the priority assigned to each network. For example, if you have two network routes, separate Double-Take and your public traffic. If you have three routes, separate the public traffic and then separate Double-Take from the cluster heartbeat.
    - **Interval to check unresponsive nodes**—Specify how much time, in seconds, between checks of nodes to see if a Double-Take connection can be made.
    - **Delay connection until resources dependent on this one are online**—This option allows you to delay a Double-Take connection until any resources that have the GeoCluster Replicated Disk resource as a dependency are online. By ensuring that all resources that are dependent on the GeoCluster Replicated disk resource are online before starting the connection, the chance of a conflict occurring because application resources are attempting to open files exclusively while Double-Take is mirroring those files is removed.
5. On the **Orphans** tab, you will see the option to delete orphan files. An orphan is a file that exists in the target location but is not in the source location. You can select to delete orphans during a mirror.
6. On the **Compression** tab, you will see your compression configuration. If you want to configure compression, verify that **Enable Compression** is selected. Depending on the compression algorithms available for your operating system, you may see a slider bar indicating different compression levels. Set the level from minimum to maximum compression to suit your needs.
7. On the **Mirror Properties** tab, specify your Double-Take mirroring settings.

- **Full Mirror**—All files will be sent from the source to the target.
- **File differences**—Only those files that are different based size or date and time will be sent from the source to the target.
  - **Send data only if Source is newer than Target**—Only those files that are newer on the source are sent to the target.

    

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

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

8. No other settings are required for the GeoCluster Replicated Disk resource, although there are optional settings available. See *GeoCluster Replicated Disk resource properties* on page 665 for details. Click **OK** to save the GeoCluster Replicated Disk configuration changes that you made.
9. To control the resource, you can bring it online and take it offline. Neither of these actions trigger failover. They just control the activity of the resource.

Creating the GeoCluster Replicated Disk Resource on Windows 2008 or 2012 Hyper-V

1. Create a virtual machine using the Hyper-V Manager. For details, see your Hyper-V documentation.
2. From the Failover Cluster Manager, cluster the virtual machine using the High Availability Wizard.
3. Take the virtual machine cluster group offline by right-clicking on it and selecting **Take this resource offline** (for Windows 2008) or **Take Offline** (for Windows 2012).
4. Right-click on the virtual machine cluster group and select select **Add a resource**, **More resources**, **Add GeoCluster Replicated Disk** (for Windows 2008) or **GeoCluster Replicated Disk** (for Windows 2012).
5. Right-click on the resource and select **Properties**.
6. On the **General** tab, specify a name that identifies which application, file set, disk, and so on that you are protecting. This name must be unique within the cluster.
7. On the **Connection Parameters** tab, specify the GeoCluster Replicated Disk connection parameters using the settings below.
   - **Disk to replicate**—Select the volume where the virtual machine .vhd or .vhdx file is stored.
   - **Network to route Double-Take mirroring and replication traffic over**—Select the network that you want to use for Double-Take mirroring and replication traffic. If you do not have multiple networks established, you will only be able to select the one network that does exist. If you do not select a network, Double-Take will use DNS to determine a network route to use. Ideally, the networks used for various traffic should be separated. This is dependent on the number of networks that you established when you created the

cluster and the priority assigned to each network. For example, if you have two network routes, separate Double-Take and your public traffic. If you have three routes, separate the public traffic and then separate Double-Take from the cluster heartbeat.

- **Interval to check unresponsive nodes**—Specify how much time, in seconds, between checks of nodes to see if a Double-Take connection can be made.
- **Delay connection until resources dependent on this one are online**—This option allows you to delay a Double-Take connection until any resources that have the GeoCluster Replicated Disk resource as a dependency are online. By ensuring that all resources that are dependent on the GeoCluster Replicated disk resource are online before starting the connection, the chance of a conflict occurring because application resources are attempting to open files exclusively while Double-Take is mirroring those files is removed.

8. On the **Orphans** tab, you will see the option to delete orphan files. An orphan is a file that exists in the target location but is not in the source location. You can select to delete orphans during a mirror.
9. On the **Compression** tab, you will see your compression configuration. If you want to configure compression, verify that **Enable Compression** is selected. Depending on the compression algorithms available for your operating system, you may see a slider bar indicating different compression levels. Set the level from minimum to maximum compression to suit your needs.
10. On the **Mirror Properties** tab, specify your Double-Take mirroring settings.
    - **Full Mirror**—All files will be sent from the source to the target.
    - **File differences**—Only those files that are different based size or date and time will be sent from the source to the target.
      - **Send data only if Source is newer than Target**—Only those files that are newer on the source are sent to the target.

        

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

      - **Use block checksum**—For those files flagged as different, the mirror performs a checksum comparison and only sends those blocks that are different.
    - **Calculate Replication Set size prior to mirror**—Determines the size of the protected data set prior to starting the mirror. The mirroring status will update the percentage complete if the data set size is calculated.
11. No other settings are required for the GeoCluster Replicated Disk resource, although there are optional settings available. See *GeoCluster Replicated Disk resource properties* on page 665 for details. Click **OK** to save the GeoCluster Replicated Disk configuration changes that you made.
12. Modify the virtual machine and virtual machine configuration resources and add the GeoCluster Replicated Disk resource (not the GeoCluster Replicated Disk Status resource) as a dependency.

Bringing the resource online

The GeoCluster Replicated Disk resource will appear offline after it is created. When you bring it online, the following actions occur.

1. A Double-Take job is created.
2. The job is connected to all of the possible owners specified in the resource (except the active node which is the source).
3. A mirror is initiated to create the baseline copy of data from the active node to all of the possible owners.
4. The drive where the mirrored data is located on each of the possible owners is made read-only to all other applications except Double-Take.
5. Real-time replication from the active node to all of the possible owners begins.

If you are using Windows 2003, right-click the resource and select **Bring online**.

If you are using Windows 2008, right-click the resource and select **Bring this resource online**.

If you are using Windows 2012, right-click the resource and select **Bring Online**.

Taking the resource offline

When you take the GeoCluster Replicated Disk resource offline, the following actions occur.

1. Real-time replication from the active node to the possible owners stops.
2. The read-only limitation is removed from the corresponding drive letters on the possible owners.
3. The job is disconnected from all of the possible owners.
4. The job is deleted.

If you are using Windows 2003, right-click the resource and select **Take offline**.

If you are using Windows 2008, right-click the resource and select **Take this resource offline**.

If you are using Windows 2012, right-click the resource and select **Take Offline**.

If the GeoCluster Replicated Disk Resource is offline, it will impact any application depending on it. Data integrity cannot be guaranteed on the other nodes in the cluster.

GeoCluster Replicated Disk resource properties

Resource properties are displayed differently in Windows 2003, Windows 2008, and Windows 2012. For example, the possible owners of a resource is listed on the **General** tab of the resource properties in Windows 2003, while in Windows 2008 they are listed on the **Advanced Policies** tab.

For both operating systems, right-click the resource and select **Properties**, when you want to view or modify the resource properties.

- See *GeoCluster Replicated Disk properties on Windows 2003* on page 665
- See *GeoCluster Replicated Disk properties on Windows 2008 or 2012* on page 668

GeoCluster Replicated Disk properties on Windows 2003

There are seven properties tabs for the GeoCluster Replicated Disk resource on Windows 2003.

1. **General**—This tab identifies the **Name** and **Description** of the resource and the **Possible Owners**. If you change the name of the resource, the replication set name will not change until the next time the resource is brought online. The GeoCluster Replicated Disk resource must have at least two possible owners to function properly. Modifying the **Possible Owners** list will cause one of the following actions to occur.
    - If you add additional **Possible Owners**, the GeoCluster Replicated Disk resource will connect the resource's replication set to the new owners and begin a mirror to each.
    - If you remove **Possible Owners**, the GeoCluster Replicated Disk resource will disconnect the resource's replication set from each owner removed.
2. **Dependencies**—By default, the GeoCluster Replicated Disk resource is not dependent on any other resources.
3. **Advanced settings**—This tab controls how and when MSCS handles a failure of the resource.
    - **Do not restart**—Select this option if you do not want cluster service to restart the resource if it fails.
    - **Restart**—Select this option if you want cluster service to restart the resource if it fails.
        - Enable **Affect the group** if you want a failure of this resource to move the group to another node. If you disable this option, cluster service still attempts to restart the resource using the **Threshold** and **Period** values, but the failure of the resource will not cause the group to move to another node.
        - The **Threshold** and **Period** values determine the number of times cluster service will attempt to restart the failed resource within a specified period of time before moving the group to another node.
    - **"Looks Alive" poll interval**—This setting specifies how often the resource is polled to determine whether it is still running on the active node. You can choose a value from the resource type, or you can specify your own value.
    - **"Is Alive" poll interval**—This setting designates how often the possible owners are polled to determine whether the specified disk on each node can be written to and read from. You can choose a value from the resource type, or you can specify your own value.
    - **Pending timeout**—This value determines how long the resource is allowed to remain in a pending state before it fails. If the resource takes longer than the time specified to be

brought online or taken offline, the resource fails.

For more information on the **Advanced Setting**s options, see your Windows documentation.

4. **Connection parameters**—This tab controls disk replication, network routing, and orphan files for GeoCluster.
   - **Disk to replicate**—The volume to replicate
   - **Network to route Double-Take mirroring and replication traffic over**—The network to use for Double-Take mirroring and replication traffic. If you do not have multiple networks established, you will only be able to select the one network that does exist. If you do not select a network, GeoCluster will use DNS to determine a network route to use.

     Ideally, the networks used for various traffic should be separated. This is dependent on the number of networks that you established when you created the cluster and the priority assigned to each network. For example, if you have two network routes, separate GeoCluster and your public traffic. If you have three routes, separate the public traffic and then separate GeoCluster from the cluster heartbeat.

     Modifications to either of the first two settings will not take effect until the next time the resource is brought online.
   - **Interval to check unresponsive nodes**—The frequency to determine how often an unresponsive node is checked to see if a Double-Take connection can be made
   - **Delay connection until resources dependent on this one are online**—This option allows you to delay a Double-Take connection until any resources that have the GeoCluster Replicated Disk resource as a dependency are online. By ensuring that all resources that are dependent on the GeoCluster Replicated disk resource are online before starting the connection, the chance of a conflict occurring because application resources are attempting to open files exclusively while GeoCluster is mirroring those files is removed.
5. **Orphans**—An orphan is a file that exists in the target location but is not in the source location. You can select to delete orphans during a mirror.
6. **Compression**—If you want to configure Double-Take compression, verify that **Enable Compression** is selected. Depending on the compression algorithms available for your operating system, you may see a slider bar indicating different compression levels. Set the level from minimum to maximum compression to suit your needs.
7. **Mirror Properties**—This tab controls the Double-Take mirroring process.
   - **Full Mirror**—All files in the replication set will be sent from the source to the target.
   - **File differences**—Only those files that are different based size or date and time will be sent from the source to the target.
     - **Send data only if Source is newer than Target**—Only those files that are newer on the source are sent to the target.

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

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

GeoCluster Replicated Disk properties on Windows 2008 or 2012

There are nine properties tabs for the GeoCluster Replicated Disk resource on Windows 2008. If you are using Windows 2012, there are the same nine properties tab, plus an additional tab which shows the private properties of the resources. These are equivalent to the properties available in the other tabs.

1. **General**—This tab identifies the **Name** and **Resource type** of the resource. It also displays the current state of the resource and an additional detailed status message.
2. **Dependencies**—By default, the GeoCluster Replicated Disk resource is not dependent on any other resources.
3. **Policies**—This tab controls how and when MSCS handles a failure of the resource. For more information on **Policies** options, see your Windows documentation.
    - **If resource fails, do no restart**—Select this option if you do not want cluster service to restart the resource if it fails.
    - **If resource fails, attempt restart on current node**—Select this option if you want cluster service to restart the resource if it fails. Specify the length of time to attempt restarts and the number of restarts to attempt during that period of time.
    - **If restart is unsuccessful, fail over all resources in this service or application**—If this option is enabled, the failure of the group will cause the resource to move to another node. If this option is disabled, the failure of the resource will not cause the resource to move to another node.
    - **If all the restart attempts fail, begin restarting again after the specified period**—If this option is enabled, the cluster will delay the length of time specified before trying to restart the resource again.
    - **Pending timeout**—This value determines how long the resource is allowed to remain in a pending state before it fails. If the resource takes longer than the time specified to be brought online or taken offline, the resource fails.
4. **Advanced Policies**—This tab controls resource specific settings. For more information on **Advanced Policies** options, see your Windows documentation.
    - **Possible owners**—All nodes of the cluster are listed. Select or deselect the nodes that you want to be possible owners.
        - If you add additional owners, the GeoCluster Replicated Disk resource will connect the resource's replication set to the new owners and begin a mirror to each.
        - If you remove owners, the GeoCluster Replicated Disk resource will disconnect the resource's replication set from each owner removed.

          The GeoCluster Replicated Disk resource must have at least two possible owners to function properly.
    - **Basic resource health check interval**—This setting is formerly known as the Looks Alive poll interval. It specifies how often the resource is polled to determine whether it is still running on the active node. You can choose the standard time period of 5 seconds, or you can specify your own value.
    - **Thorough resource health check interval**—This setting is formerly known as the Is Alive poll interval. It designates how often the possible owners are polled to determine whether the specified disk on each node can be written to and read from. You can choose the standard time period of 1 minute, or you can specify your own value.

- **Run this resource in a separate Resource Monitor**—You should enable this option so that each GeoCluster Replicated Disk resource runs in its own monitor.

5. **Connection parameters**—This tab controls disk replication, network routing, and orphan files for Double-Take.
   - **Disk to replicate**—The volume to replicate
   - **Network to route Double-Take mirroring and replication traffic over**—The network to use for Double-Take mirroring and replication traffic. If you do not have multiple networks established, you will only be able to select the one network that does exist. If you do not select a network, GeoCluster will use DNS to determine a network route to use.

     Ideally, the networks used for various traffic should be separated. This is dependent on the number of networks that you established when you created the cluster and the priority assigned to each network. For example, if you have two network routes, separate GeoCluster and your public traffic. If you have three routes, separate the public traffic and then separate GeoCluster from the cluster heartbeat.

     Modifications to either of the first two settings will not take effect until the next time the resource is brought online.
   - **Interval to check unresponsive nodes**—The frequency to determine how often an unresponsive node is checked to see if a Double-Take connection can be made
   - **Delay connection until resources dependent on this one are online**—This option allows you to delay a Double-Take connection until any resources that have the GeoCluster Replicated Disk resource as a dependency are online. By ensuring that all resources that are dependent on the GeoCluster Replicated disk resource are online before starting the connection, the chance of a conflict occurring because application resources are attempting to open files exclusively while GeoCluster is mirroring those files is removed.
6. **Orphans**—An orphan is a file that exists in the target location but is not in the source location. You can select to delete orphans during a mirror.
7. **Compression**—If you want to configure Double-Take compression, verify that **Enable Compression** is selected. Depending on the compression algorithms available for your operating system, you may see a slider bar indicating different compression levels. Set the level from minimum to maximum compression to suit your needs.
8. **Online Pending**—Because context-sensitive, right-click menus are not available in the Windows 2008 Failover Cluster Administrator, GeoCluster processing controls have been added to a properties tab. For details on this tab, see *Monitoring and controlling GeoCluster jobs* on page 671.
9. **Mirror Properties**—This tab controls the Double-Take mirroring process.
   - **Full Mirror**—All files in the replication set will be sent from the source to the target.
   - **File differences**—Only those files that are different based size or date and time will be sent from the source to the target.
     - **Send data only if Source is newer than Target**—Only those files that are newer on the source are sent to the target.

       

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

- **Use block checksum**—For those files flagged as different, the mirror performs a checksum comparison and only sends those blocks that are different.

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

Monitoring and controlling GeoCluster jobs

The Double-Take job created by the GeoCluster will be displayed as an **Legacy Job** on the **Manage Jobs** page of the Double-Take Console. You will have limited control over the job, including taking and managing snapshots, viewing the job log, initiating mirrors and verifications, setting bandwidth limitations, deleting orphan files, and pausing and resuming the target.

If you delete the job, the GeoCluster Replicated Disk resource will not be deleted. You will have to remove and readd the passive nodes to the resource to re-create the job without impacting production. If you are not worried about impacting production, you can take the resource offline and then bring it back online to re-create the job.

Ideally, you should use the standard Windows cluster tools to monitor the status of the resource. See your cluster documentation for details on monitoring a cluster resource.

You can also use the following information to help you monitor the GeoCluster Replicated Disk resource.

- Do not use the **Initiate Failure** feature of Cluster Administrator to test failover of GeoCluster resources. Use other test methods, such as manually moving the group or unplugging the owning nodes network cable.
- Do not use the **Automatic Failback** feature of Cluster Administrator. If you need to return ownership to the original node, wait until GeoCluster has completed mirroring from the new owning node back to the original owning node and then manually move the group.
- If you change an IP address on any node of the cluster, you must stop and restart the cluster service on all of the nodes in the cluster in order for GeoCluster to detect the new IP address.
- If you must reboot the owning node, you should move all of your cluster groups to another node before the reboot. This will ensure data integrity and allow you to make sure the applications come online before the node is rebooted.

Resolving an online pending GeoCluster Replicated Disk resource

When the GeoCluster Replicated Disk resource is in an online pending state, you are protected from possible data corruption. If you are using Windows 2003, review the description of the GeoCluster Replicated Disk Status resource to see why the GeoCluster Replicated Disk resource is in the online pending state. If you are using Windows 2008 or 2012, you can see the online pending status directly in the description of the GeoCluster Replicated Disk resource. If the pending state were bypassed, the node where you are trying to bring the resource online would have incomplete data, which would then be replicated to the other nodes in the cluster. This state safeguards you from corrupting your data.

There are different options for resolving an online pending state, depending on whether your operating system supports snapshots. Therefore, some of the following options may not be displayed or may be disabled if they are not valid for your configuration.

If you are using Windows 2003, right-click on the online pending resource and select the desired control. The controls are described in the following tables.

If you are using Windows 2008 or 2012, right-click the online pending resource, select **Properties**, select the **Online Pending** tab, and click the desired control. The controls are described in the following tables.

- **Revert**
  - **Windows 2003 menu**—Revert to snapshot
  - **Windows 2008 or 2012 menu**—Revert Snapshot
  - **Description**—If you have a snapshot of the target data available, you can revert to that data. If you revert to a snapshot, any data changes made after the snapshot's specified date and time will be lost. A Double-Take connection will be established to replicate the node's data (at the snapshot point-in-time) to the other nodes.
- **Discard queue**
  - **Windows 2003 menu**—Discard target queue
  - **Windows 2008 or 2012 menu**—Discard Queue
  - **Description**—If you have data in the target queue, you can discard that data. If you discard the queued data, you will lose the changes associated with that data made on the previously owning node. A Double-Take connection will be established to replicate the node's data (without the data that was in queue) to the other nodes.
- **Offline or fail**
  - **Windows 2003 menu**—Force Resource Offline
  - **Windows 2008 or 2012 menu**—Fail Resource
  - **Description**—If you are using Windows 2003, you can force the resource offline. If you are using Windows 2008 or 2012, you can fail the resource. In either case, no Double-Take connection will be established.
- **Verify group**
  - **Windows 2003 menu**—Verify Group
  - **Windows 2008 or 2012 menu**—Verify Group
  - **Description**—With this option and snapshot capability, you can test the data on the node before deciding whether to use it. If you select this option, a snapshot of the node's current

Double-Take data will be taken, the disk will come online, but the GeoCluster Replicated Disk resource will not come online, allowing you to check the data. (This means there is no Double-Take connection established at this time.) Once the snapshot is taken, you can test the data on the node to see if it is viable. Make sure you prevent user access while you are verifying the data. Once you have tested the data, you need to right-click on the online pending resource again and accept or reject the data.

- **Accept data**
  - **Windows 2003 menu**—Accept Data
  - **Windows 2008 or 2012 menu**—Accept
  - **Description**—If you accept the data, the current data on the node will be used, and a Double-Take connection will be established to replicate the current node's data to the other nodes. If any other nodes in the cluster contain more recent data, this node will overwrite that data and it will be lost.
- **Reject data**
  - **Windows 2003 menu**—Reject Data
  - **Windows 2008 or 2012 menu**—Reject
  - **Description**—If you reject the data, the node will be reverted to the snapshot that was taken when you selected the **Verify Group** option. Any changes made on the node after that snapshot was created will be lost. This option essentially takes you back to where you were, allowing you the opportunity to check other nodes for more recent data.

    If you have multiple GeoCluster Replicated Disks in the same group and have selected to reject the data after verifying the group, the rejection processing may take several minutes.

GeoCluster Replicated Disk Status Resource

The function of the GeoCluster Replicated Disk Status resource (also displayed as GRD Status) varies between Windows 2003, 2008, and 2012. In both operating systems, it is automatically created when the first GeoCluster Replicated Disk resource is created in a group. Once the status resource is created, it will exist as long as there is a GeoCluster Replicated Disk resource in the group. When the last GeoCluster Replicated Disk resource in a group is deleted, the status resource will be deleted. Only one status resource is created per group.If the resource is deleted, it will automatically be re-created.

If you are using Windows 2003, the description of the status resource corresponds to various states of your Double-Take data. By reviewing the status descriptions, you can tell at-a-glance the state of your Double-Take data. If you are using Windows 2008 or 2012, these status descriptions are seen directly in the GeoCluster Replicated Disk resource description, rather than the status resource.

For example, you may see the status "The status of all targets is OK." This indicates the data on each target node is in a good state. Another message may be "Target target_name is queuing. Data in queue on target." This indicates the data on the specified target is not up-to-date. Because there is data in queue on the target, that has not been written to disk yet, the target data is out-of-date. Or you may see either of the following status descriptions.

- Target target_name is pending. Data integrity not guaranteed.
- Target target_name is suspect. Data integrity not guaranteed.

These messages indicate the data on the specified target node is not in a good state. This could be because a mirror is in progress, an operation has been dropped on the target, or another Double-Take processing issue.As long as the status is pending, data integrity cannot be guaranteed on the specified target node. Check the Double-Take logs for more information. See *Log files* on page 677 for more details.

The text of the descriptions may vary between Windows 2003, 2008, and 2012.

Another function of the status resource, for all Windows versions, is to keep you from moving the GeoCluster Replicated Disk resource to another node at the wrong time and potentially corrupting your data. If the GeoCluster Replicated Disk resource was moved while the status resource is in a pending or queuing state, the new node would have incomplete data, which would then be replicated to the other nodes in the cluster. This resource safeguards you from corrupting your data. This happens by removing passive nodes as possible owners and discarding any manual changes made to the possible owners list.

Chapter 17 Simulating protection

Double-Take offers a simple way for you to simulate protection 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 DTInfo utility 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. The DTInfo utility can be found on the product DVD, in the Double-Take installation directory, or on the Vision Solutions support web site.

1. From the source where you will be running the TDU, run DTInfo.exe. It may take several minutes for DTInfo to finish processing. After DTInfo processing is complete, a \support subdirectory will automatically be created in the Double-Take installation directory. (The default installation directory is \Program Files\Vision Solutions\Double-Take.) A .zip file will contain the information gathered from DTInfo. The file name is based on the machine name. To distinguish this file from the next time you run DTInfo, append a unique identifier, perhaps the date and time, to the end of the file name.
2. Establish a protection job, noting the following caveats.
    - The TDU is not available for the following job types.
        - Full server to ESX appliance
        - V to ESX
        - V to Hyper-V
        - Agentless vSphere
    - When you get to the **Choose Target Server** page in the workflow, select the **Diagnostics** target.
    - 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, if you have selected a job type for an ESX server, you will not need to specify options for the target replica virtual machine 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 *Statistics* on page 688.
4. After your simulation is complete, repeat step 1 to run DTInfo again, appending the new unique identifier to the end of the new .zip file.

Chapter 18 Monitoring tools

Outside of the Double-Take consoles, you have other general monitoring tools available for all job types.

Log files

Double-Take generates log files to gather alerts, which are notification, warning, and error messages.

- **Double-Take log**—This log records data from the Double-Take service, also referred to as the Double-Take engine. The Double-Take service controls the data movement functions like Double-Take mirroring and replication. This log file can be viewed from within the Double-Take Console (see *Viewing the log files through the Double-Take Console* on page 678) or through any standard text editor (see *Viewing the log files through a text editor* on page 682). You can also filter the data in the log file using the Double-Take LogViewer utility. (see *Filtering the log file with LogViewer* on page 686).
- **Double-Take Management Service log**—This log records data from the Double-Take Management Service. It controls all non-data movement aspects of each job. This log file can be viewed from within the Double-Take Console (see *Viewing the log files through the Double-Take Console* on page 678) or through any standard text editor (see *Viewing the log files through a text editor* on page 682).
- **Double-Take job log**—This log records job specific message. There is a unique job log for each job you create. This log file can be viewed from within the Double-Take Console (see *Viewing the log files through the Double-Take Console* on page 678) or through any standard text editor (see *Viewing the log files through a text editor* on page 682).
- **Double-Take Console log**—This log records data and user interaction from the Double-Take Console. This log file can be viewed through any standard text editor (see *Viewing the log files through a text editor* on page 682).
- **Virtual Recovery Appliance log**—For full server to ESX appliance jobs, there are log files maintained on the appliance that record appliance processing. These log files are called master.log and job-<uid>.log, where uid is a unique identifier. You can find them on the appliance under /opt/dbtk/log. You can use any standard text editor to view the log files. If you are looking for source specific logging, check the log file on the source which is the /var/log/messages file.
- **Controller and Replication Appliance log**—Similar to the engine log, the controller and replication appliances have a log file to track appliance processing and the data movement functions of mirroring and replication. These log files can be viewed from with in the Double-Take Console (see *Viewing the log files through the Double-Take Console* on page 678). You can also view the file directly. The log for the controller appliance is /opt/visionsolutions/doubletake/log/controller-service.log. The log for the replication appliance is /opt/visionsolutions/doubletake/log/management-service.log.

For cluster environments, be sure and review the log files on all nodes of the cluster.

Viewing the log files through the Double-Take Console

You can view the Double-Take, Double-Take Management Service, controller and replication appliances, job, and Double-Take Console log files through the Double-Take Console.

Viewing the Double-Take, Double-Take Management Service, and controller and replication appliance logs

You can view the Double-Take and Double-Take Management Service log files through the Double-Take Console using either of these two methods.

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

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

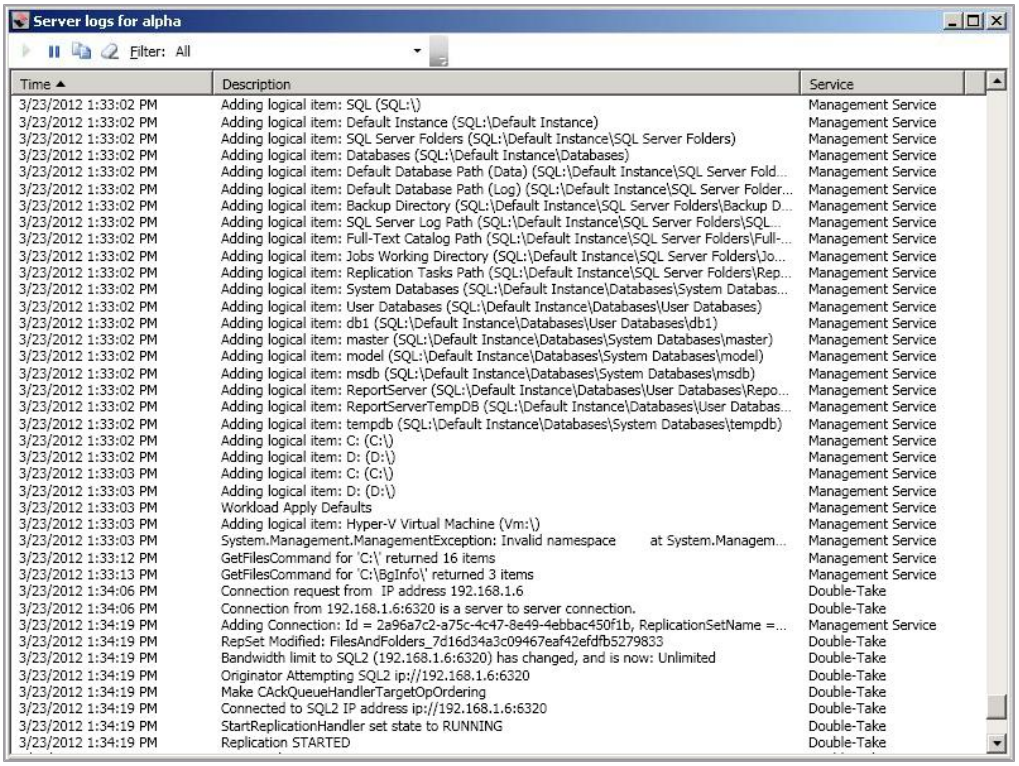

The following table identifies the controls and the table columns in the **Server logs** window.

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.

Viewing the job log file

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

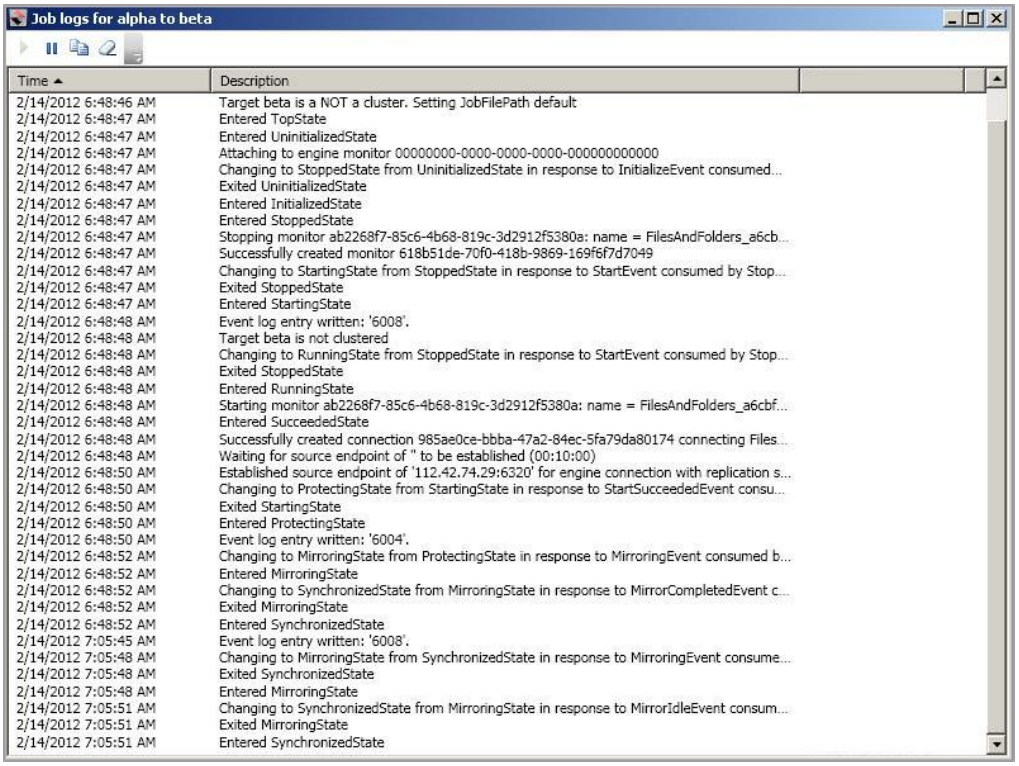

The following table identifies the controls and the table columns in the **Job logs** window.

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.

Viewing the Double-Take Console log file

You can view the Double-Take Console log file by using the following instructions.

1. Select **Tools**, **Options**.
2. Expand the **Diagnostics** section, if necessary.
3. Click **View Log File**.
4. Click **Cancel** to return back to the console page you were on previously. (Unless you were in the middle of a job creation workflow, in which case you will be returned to the beginning of the workflow.)

The log file is opened in the default text editor. The file will remain open until you close.

Viewing the log files through a text editor

You can view the Double-Take, Double-Take Management Service, job, and Double-Take Console log files through any text editor.

Viewing the Double-Take log

This log file is located in the same directory where you installed Double-Take. By default, this is \Program Files\Vision Solutions\Double-Take.

The Double-Take log file consists of a base name, a series number, and an extension. The base name is dtlog and the extension is .dtl. The series number ranges from 1 to 999. For example, Double-Take begins logging messages to dtlog1.dtl. When this file reaches its maximum size, which by default is 5 MB, the next log file will be written to dtlog2.dtl. As long as log messages continue to be written, files dtlog3.dtl, dtlog4.dtl, and dtlog5.dtl will be opened and filled. When the maximum number of files is reached, which by default is 5, the oldest file is deleted. For example, when dtlog6.dtl is created, dtlog1.dtl is deleted, and when dtlog7.dtl is created, dtlog2.dtl is deleted. When file dtlog999.dtl is created and filled, dtlog1.dtl will be re-created and Double-Take will continue writing log messages to that file. The Double-Take log file settings can be modified for each server. See *Log file properties* on page 100 for details.

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

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

```
02/07/2013 05:26:32.954 1360 3668 1 2 0 Using default heap
02/07/2013 05:26:33.048 1360 3668 3 2 0 Winsock v2.2Enabled
02/07/2013 05:26:35.689 1360 3668 5 2 0 Virtual operating system detected.
02/07/2013 05:26:35.689 1360 3668 6 2 69 Kernel Starting on BETA ip://112.42.74.30:6320 V
02/07/2013 05:26:35.689 1360 3668 7 2 0 WARNING: Unable to determine Firewall status, ass
02/07/2013 05:26:35.751 1360 3668 8 2 69 Kernel Started on BETA ip://112.42.74.30:6320 Ve
02/07/2013 05:26:35.954 1360 3668 9 2 503008 LoadNL set state to IDLE
02/07/2013 05:26:35.954 1360 3944 10 2 503008 Op Retrieval: entering standard retrieval m
02/07/2013 05:26:35.954 1360 3020 11 2 503008 Op Parsing: entering standard parsing mode
02/07/2013 05:26:36.017 1360 3668 12 2 0 The local security process has been added to the
02/07/2013 05:26:38.751 1360 3668 14 2 0 SourceRelationships::UpdateSourceMap:GUID(A0779C
02/07/2013 05:26:38.751 1360 3668 15 2 0 Source IP '10.10.10.29:6320' for ServerID 'GUID(
02/07/2013 05:26:38.814 1360 3668 16 2 52000 New Target Internals created for 10.10.10.29
02/07/2013 05:26:38.876 1360 3668 17 2 302004 Queuing to disk has started
02/07/2013 05:26:38.876 1360 3668 18 2 302005 Disk Queue is empty - Queuing to disk has s
02/07/2013 05:26:38.939 1360 3668 19 2 52501 Target module loaded successfully
02/07/2013 05:26:39.064 1360 3816 20 2 71 Originator Attempting ip://10.10.10.29:6320
02/07/2013 05:26:39.157 1360 2260 21 2 72 Connection request from IP address 112.42.74.30
```

Viewing the Double-Take Management Service log

This file is located in the \Service\Logs subdirectory where you installed Double-Take.

The Double-Take Management Service log file consists of a base name, an optional date, an optional series number, and an extension. The base name is ManagementService and the extension is .log. When this file reaches its maximum size, 10 MB, the file will be renamed with the current date in year, month, day format. For example, it might be ManagementService.20130207.log. The latest log messages are then stored in ManagementService.log. If the main file fills again on the same day, then a series number will be used. In this case, the ManagementService.log file will be renamed to ManagementService.20130207.1.log. If the main file is filled on a different day, that date will be specified. In each case, the latest log messages will be stored in ManagementService.log. When the maximum number of files is reached, which is 5, the oldest file is deleted. The Management Service log file settings cannot be modified.

```
[2013-02-07 05:27:19] Verbose: Successfully opened service host JobManager (PID:2512, TID
[2013-02-07 05:27:19] Verbose: Successfully opened service host WorkloadManager (PID:2512
[2013-02-07 05:27:19] Verbose: Successfully opened service host ManagementService (PID:25
[2013-02-07 05:27:19] Verbose: Successfully opened service host CoreEngine (PID:2512, TID
[2013-02-07 05:27:19] Verbose: Successfully opened service host CoreSystemState (PID:2512
[2013-02-07 05:27:19] Verbose: Successfully opened service host DiskManager (PID:2512, TI
[2013-02-07 05:27:20] Verbose: Completed: 10.10.10.30 (PID:2512, TID:9, AID:fce50580-0307
[2013-02-07 05:27:20] Information: The Double-Take engine is initialized. (PID:2512, TID:
[2013-02-07 05:27:20] Verbose: The Double-Take engine source module is initialized. (PID:
[2013-02-07 05:27:20] Verbose: The Double-Take engine target module is initialized. (PID:
[2013-02-07 05:27:20] Verbose: Removing Connection: Id = f34dcb1e-ac6b-4267-a891-5e45629b
[2013-02-07 05:27:20] Verbose: The Double-Take engine failover module is initialized. (PI
[2013-02-07 05:27:21] Verbose: Removing Monitor: Id = 5fa3b4eb-ea21-4e47-976e-25fd62c10c5
[2013-02-07 06:48:44] Verbose: Entered TopState (PID:2512, TID:13, AID:7c5f8158-16f0-4330
[2013-02-07 06:48:44] Verbose: Entered UninitializedState (PID:2512, TID:13, AID:7c5f8158
[2013-02-07 06:48:44] Verbose: Changing to StoppedState from UninitializedState in respon
[2013-02-07 06:48:44] Verbose: Exited UninitializedState (PID:2512, TID:13, AID:7c5f8158-
[2013-02-07 06:48:44] Verbose: Entered InitializedState (PID:2512, TID:13, AID:7c5f8158-1
```

Viewing the job log

This file is located on the target server in the \Service\Logs subdirectory where you installed Double-Take.

The job log file consists of a global unique identifier (GUID) for the job and the job name. When this file reaches its maximum size, 10 MB, the file will be renamed with the current date in year, month, day format. For example, it might be a6cbf990-ba67-403d-855f-5bb44c18e1d6 (alpha to beta).20130207.log. The latest log messages are then stored in the base GUID_name.log file. If the main file fills again on the same day, then a series number will be used. In this case, the example file would be renamed to a6cbf990-ba67-403d-855f-5bb44c18e1d6 (alpha to beta).20130207.1.log. If the main file is filled on a different day, that date will be specified. In each case, the latest log messages will be stored in the main GUID_name.log file. When the maximum number of files is reached, which is 5, the oldest file is deleted. The job log file settings cannot be modified.

```
[2013-02-07 06:48:46] Verbose: Entered TopState (PID:2512, TID:13, AID:e6246673-57a9-4a25
[2013-02-07 06:48:46] Verbose: Entered UninitializedState (PID:2512, TID:13, AID:e6246673
[2013-02-07 06:48:46] Information: Target beta is a NOT a cluster. Setting JobFilePath de
[2013-02-07 06:48:47] Verbose: Entered TopState (PID:2512, TID:13, AID:e6246673-57a9-4a25
[2013-02-07 06:48:47] Verbose: Entered UninitializedState (PID:2512, TID:13, AID:e6246673
[2013-02-07 06:48:47] Verbose: Attaching to engine monitor 00000000-0000-0000-0000-000000
[2013-02-07 06:48:47] Verbose: Changing to StoppedState from UninitializedState in respon
[2013-02-07 06:48:47] Verbose: Exited UninitializedState (PID:2512, TID:5, AID:e6246673-5
[2013-02-07 06:48:47] Verbose: Entered InitializedState (PID:2512, TID:5, AID:e6246673-57
[2013-02-07 06:48:47] Verbose: Entered StoppedState (PID:2512, TID:5, AID:e6246673-57a9-4
[2013-02-07 06:48:47] Verbose: Stopping monitor ab2268f7-85c6-4b68-819c-3d2912f5380a: nam
[2013-02-07 06:48:47] Information: Successfully created monitor 618b51de-70f0-418b-9869-1
[2013-02-07 06:48:47] Verbose: Changing to StoppedState from UninitializedState in respon
[2013-02-07 06:48:47] Verbose: Exited UninitializedState (PID:2512, TID:13, AID:e6246673-
[2013-02-07 06:48:47] Verbose: Entered InitializedState (PID:2512, TID:13, AID:e6246673-5
[2013-02-07 06:48:47] Verbose: Entered StoppedState (PID:2512, TID:13, AID:e6246673-57a9-
[2013-02-07 06:48:47] Verbose: Changing to StartingState from StoppedState in response to
[2013-02-07 06:48:47] Verbose: Exited StoppedState (PID:2512, TID:13, AID:5decb7a7-ff7a-4
[2013-02-07 06:48:47] Verbose: Entered StartingState (PID:2512, TID:13, AID:5decb7a7-ff7a
[2013-02-07 06:48:48] Information: Event log entry written: '6008'. (PID:2512, TID:13, AI
[2013-02-07 06:48:48] Verbose: Target beta is not clustered (PID:2512, TID:13, AID:5decb7
[2013-02-07 06:48:48] Verbose: Changing to RunningState from StoppedState in response to
[2013-02-07 06:48:48] Verbose: Exited StoppedState (PID:2512, TID:14, AID:e6246673-57a9-4
[2013-02-07 06:48:48] Verbose: Entered RunningState (PID:2512, TID:14, AID:e6246673-57a9-
[2013-02-07 06:48:48] Verbose: Starting monitor ab2268f7-85c6-4b68-819c-3d2912f5380a: nam
```

Viewing the Double-Take Console log

This file is called Double-Take Console.log and its location depends on your operating system. If you are using Windows 2008 or 2012, the file will be located in \Users\<your user name>\AppData\Local\Vision Solutions\Double-Take Console\Logs. If you are using Windows 2003, the log file will be located in \Documents and Settings\<your user name>\Local Settings\Application Data\Vision Solutions\Double-Take Console\Logs. Note that these are hidden locations, so you will have to search directly for the file name or display hidden files to browse for it. Also note that the log file is dependent on the user who is logged in, so there will be multiple Double-Take Console log files if you have multiple Double-Take users.

```
[2013-02-07 05:28:01] Verbose: Activating DoubleTake.Virtualization.VRA.Console.VRAConsol
[2013-02-07 05:28:01] Verbose: Activating DoubleTake.Virtualization.VRA.Console.VRABindin
[2013-02-07 05:28:01] Verbose: Activating DoubleTake.Virtualization.VRA.Console.Model.VRA
[2013-02-07 05:28:02] Verbose: Rendering Tier = 0 (PID:620, TID:1, AID:00000000-0000-0000
[2013-02-07 05:28:02] Verbose: Component Versions:
        Microsoft .NET Framework: 2.0.50727.3053
        Double-Take Console Framework: 7.0.0.1124
        Double-Take Applications Console Library: 7.0.0.1124
        Double-Take Dashboard Console Library: 7.0.0.1124
        Double-Take Diagnostics Console Library: 7.0.0.1124
        Double-Take Full Server Console Library: 7.0.0.1124
        Double-Take Move Console Library: 7.0.0.1124
        Double-Take for Hyper-V Console Library: 7.0.0.1124
        Double-Take Virtualization UVRA Console Library: 7.0.0.1124
        Double-Take Virtualization VRA Console Library: 7.0.0.1124
     (PID:620, TID:1, AID:2e1d2989-7b38-4c9a-93b9-063a3c12b6f3 - Extensibility)
[2013-02-07 05:29:33] Information: Qualify truealpha (PID:620, TID:1, AID:00000000-0000-0
[2013-02-07 05:40:52] Information: Qualify truealpha (PID:620, TID:1, AID:00000000-0000-0
[2013-02-07 06:06:38] Information: Qualify truealpha (PID:620, TID:1, AID:00000000-0000-0
[2013-02-07 06:25:02] Information: Qualify truealpha (PID:620, TID:1, AID:00000000-0000-0
[2013-02-07 06:47:48] Information: Qualify truealpha (PID:620, TID:1, AID:00000000-0000-0
[2013-02-07 06:48:46] Verbose: Creating job with options:
    <JobOptions xmlns:i="http://www.w3.org/2001/XMLSchema-instance" xmlns="http://schemas
      <ApplicationOptions i:nil="true" />
      <BandwidthOptions xmlns:d2p1="http://schemas.datacontract.org/2004/07/DoubleTake.Jo
        <d2p1:Entries i:nil="true" />
        <d2p1:Limit>0</d2p1:Limit>
        <d2p1:Mode>NotLimited</d2p1:Mode>
        <d2p1:Specifications xmlns:d3p1="http://schemas.datacontract.org/2004/07/DoubleTa
      </BandwidthOptions>
      <ClusterOptions>
```

Filtering the log file with LogViewer

You can filter the Double-Take log file through the Double-Take LogViewer utility. From a command prompt, use the LogViewer command from the directory where Double-Take is installed. Press Ctrl-C to exit back to the command prompt.

Command

LOGVIEWER

Description

The Double-Take logging utility that filtersDouble-Take log files

Syntax

LOGVIEWER [-PATH <***path***>] [-TYPE <***number***>] [-INCLUDE <***list***>] [-EXCLUDE <***list***>] [-NODATE] [-NOTIME] [-NOPID] [-NOTID] [-NOSEQ] [-NOTYPE] [-NOID] [-HELP]

Options

- PATH ***path***—Specify the full path to the log file
- TYPE ***number***—Allows you to filter the messages that are displayed. Specify 1 to display warning and error messages or specify 2 to display warnings, errors, and information messages.
- INCLUDE—Only includes specified IDs. All other IDs will not be displayed in the output
- EXCLUDE—Excludes specified IDs. Ignore the specified IDs and display all others
- ***list***—A comma-separated list of IDs or ID ranges that follows the INCLUDE and EXCLUDE switches. A space should separate the switch from the list but within the list, there should be no spaces. Ranges are specified with a begin and end number and separated with a dash (-).
- NODATE—Does not display the date in the output
- NOTIME—Does not display the time in the output
- NOPID—Does not display the process ID in the output
- NOTID—Does not display the thread ID in the output
- NOSEQ—Does not display the sequence number in the output
- NOTYPE—Does not display the message type number in the output
- NOID—Does not display the LogViewer ID in the output
- HELP—Displays the command options

Examples

- LogViewer -type 2
- LogViewer -include 200,400-500,10000-15000

Notes

The default setting is -type 2 which displays both type 1 and 2 messages.

Statistics

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

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

Sample DTStat output

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

Viewing the statistics file

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

Command

DTSTAT

Description

Starts the DTStats statistics logging utility from a command prompt

Syntax

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

Options

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

Examples

- DTStat -p -f statistic.sts -s statistic.csv
- DTStat -p -f statistic.sts -st statistic.csv

Notes

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

Statistics

The following table identifies the Double-Take statistics.

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

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

Statistic values are cumulative. For example if Kernel, RepBytesGenerated is 10000 at 1:00pm and 25000 at 2:00pm, the difference is 15000 and that is the amount of change that occurred within that one hour.

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

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

Date/Time Stamp

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

System Allocator, Total Bytes

The number of bytes currently allocated to the system pagefile

IQAllocator, Total Bytes

The number of bytes currently allocated to the intermediate queue

Security, Logins

The number of successful login attempts

Security, Failed Logins

The number of failed login attempts

Kernel, SourceState

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

- 2—Source is running with the replication driver

Kernel, TargetState

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

Kernel, Start Time

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

Kernel, RepOpsGenerated

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

Kernel, RepBytesGenerated

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

Kernel, MirOpsGenerated

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

Kernel, MirBytesGenerated

The number of mirror bytes transmitted to the target. This is the number of bytes generated during mirroring. In other words, this is roughly the amount of traffic being sent across the network that is generated by the mirror. It does not take into account TCP/IP overhead (headers and such), however it does account for attributes and other overhead associated with creating a file. With many small files in a directory, you will see larger statistics than expected because of the file creation overhead. Any subsequent remirror will reset this field to zero and increment from there.

Kernel, FailedMirrorCount

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

Kernel, FailedRepCount

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

Kernel, ActFailCount

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

incorrect codes, etc.

Kernel, TargetOpenHandles

The number of handles currently open on the target

Kernel, DriverQueuePercent

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

Target, PeerAddress

The IP address of the source machine

Target, LocalAddress

The IP address of the target machine.

Target, Ops Received

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

Target, Mirror Ops Received

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

Target, Retries

The number of retries performed before all operations were completed

Target, OpsDropped

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

Target, Ops Remaining

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

Target, Orphan Files Removed

The number of orphan files removed from the target machine

Target, Orphan Directories Removed

The number of orphan directories removed from the target machine

Target, Orphan Bytes Removed

The number of orphan bytes removed from the target machine

Target, Bytes In Target Queue

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

Target. Bytes In Target Disk Queue

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

Target, TasksSucceeded

The number of task commands that have succeeded on the target

Target, TasksFailed

The number of task commands that have failed on the target

Target, TasksIgnored

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

Source, autoDisConnects

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

Source, autoReConnects

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

Source, lastFileTouched

The last filename that had a replication operation executed

Connection, conPeerAddress

The IP address of the target machine

Connection, connectTime

The time that this connection was established

Connection, conState

The state of the active connection

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

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

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

Connection, conOpsInCmdQueue

The number of operations waiting to be executed on the target

Connection, conOpsInAckQueue

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

Connection, conOpsInRepQueue

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

Connection, conOpsInMirQueue

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

Connection, conBytesInRepQueue

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

Connection, conOpsTx

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

Connection, conBytesInMirQueue

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

Connection, conBytesTx

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

Connection, conBytesCompressedTx

The number of compressed bytes transmitted to the target.

Connection, conOpsRx

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

Connection, conBytesRx

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

Connection, conResentOpCount

The number of operations resent because they were not acknowledged

Connection, conBytesInDiskQueue

The number of bytes in the source disk queue

Connection, conBandwidthLimit

The amount of bandwidth that may be used to transfer data

Connection, conBytesSkipped

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

Connection, conMirrorBytesRemaining

The number of mirror bytes remaining to be transmitted

Connection, conMirrorPercent

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

Connection, conTaskCmdsSubmitted

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

Connection, conTaskCmdsQueued

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

Connection, conTasksSucceeded

The number of task commands that have succeeded on the source

Connection, conTasksFailed

The number of task commands that have failed on the source

Connection, conTasksIgnored

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

Replication service view

You can view the replication service details for a server by right-clicking on a server on the **Manage Servers** page and selecting **View Replication Service Details**. A separate window will open allowing 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 by using the drop-down list of servers in the toolbar.

The left pane of the **Replication service view** window is divided into the root and three folders.

- **Root**—This section shows high-level overview information for the server. See *Root items* on page 698 for the items in this **Server Properties** section.
- **Connections**—This section shows any active connections from this server. A connection is the underlying component of a job that controls data movement, like mirroring and replication. The Double-Take engine controls the connection.

  A connection may or may not be associated with a job. If it is not associated with a job, it can be deleted. However, you should be certain it is not associated with a job because deleting a connection that is being used can corrupt its parent job. Use the **Delete** button in the toolbar to delete a connection. (You cannot delete a GeoCluster connection.)

  When you highlight the **Connections** folder in the left pane, all active connections from this server will be displayed in the right pane. See *Connections folder items* on page 698 for details on the data displayed in this view. If you highlight a specific connection under the **Connections** folder, only the information for that connection will be displayed in the right pane. The connections are identified by the type of job and the connection ID. See *Specific connection items* on page 702 for details on the data displayed in this view.

- **Replication sets**—This section shows any replication sets on this server. The replication set is the data that your job is protecting.

  A replication set may or may not be associated with a connection. If it is not associated with a connection, it can be deleted. Use the **Delete** button in the toolbar to delete a replication set.

  When you highlight the **Replication sets** folder in the left pane, all replication sets on this server will be displayed in the right pane. See *Replication sets folder items* on page 706 for details on the data displayed in this view. If you highlight a specific replication set under the **Replication sets** folder, only the information for that replication set will be displayed in the right pane. See *Specific replication set items* on page 707 for details on the data displayed in this view.

- **Target connection entries**—This section is like the **Connections** section, however it shows any active connections to this server.

  You may see target connections that are not associated with a job. These connections will have a **Target Data State** of **Disconnected**. For example, this may happen if you delete a stopped job. These disconnected target connections cannot be deleted. They will be reused when a new connection (same job type, same servers) is created.

  When you highlight the **Target connection entries** folder in the left pane, all active connection to this server will be displayed in the right pane. See *Target connection entries folder items* on page 707 for details on the data displayed in this view. If you highlight a specific connection under the **Target connection entries** folder, only the information for that connection will be displayed in

the right pane. See *Specific target connection items* on page 709 for details on the data displayed in this view.

Root items

Server name

The name of the server

Product version

The Double-Take version

Operating system

The operating system version and edition

Source module

Indicates if the source module is running on the server

Target module

Indicates if the target module is running on the server

Failover module

Indicates if the failover module is running on the server

Connections folder items

Replication Set

The name of the replication set the connection is using

Connection ID

The incremental counter used to number connections. The number is incremented when a connection is created. It is also incremented by internal actions, such as an auto-disconnect and auto-reconnect. The lowest available number (as connections are created, stopped, deleted, and so on) will always be used. The counter is reset to one each time the Double-Take service is restarted.

Target Name

The name of the target the connection is using, including the port number

Target IP

The target IP address and port, as well as the location on the target where the replication set data is being stored. This is sometimes called the transform path.

Target Data State

- **OK**—The data on the target is in a good state.
- **Mirroring**—The target is in the middle of a mirror process. The data will not be in a good state until the mirror is complete.
- **Mirror Required**—The data on the target is not in a good state because a remirror is required. This may be caused by an incomplete or stopped mirror or an operation may have been dropped on the target.
- **Restore Required**—The data on the source and target do not match because of a failover condition. Restore the data from the target back to the source. If you want to discard the changes on the target, you can remirror to resynchronize the source and target.
- **Snapshot Reverted**—The data on the source and target do not match because a snapshot has been applied on the target. Restore the data from the target back to the source. If you want to discard the changes on the target, you can remirror to resynchronize the source and target.
- **Busy**—The source is low on memory causing a delay in getting the state of the data on the target.
- **Not Loaded**—Double-Take target functionality is not loaded on the target server. This may be caused by an activation code error.
- **Not Ready**—For a full server to ESX appliance job, the Linux drivers have not yet completed loading on the target.
- **Unknown**—The console cannot determine the status.

This field is not applicable to agentless vSphere jobs.

Target Status

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

Transmit Mode

Different jobs have different transmit mode types. The following list shows all possible transmit mode types.

- **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.

Mirror Status

Different jobs have different mirror status types. The following list shows all possible mirror status types.

- **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.
- **Restoring**—Data is being restored from the target to the source.
- **Unknown**—The console cannot determine the status.

Replication Status

Different jobs have different replication status types. The following list shows all possible replication status types.

- **Replicating**—Data is being replicated to the target.
- **In Progress**—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.

Queued (ops)

The total number of mirror and replication operations that are in the source queue

Sent (bytes)

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

Sent Compressed (bytes)

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**.

Intermediate Queue (bytes)

The total amount of memory being used by the operations buffer queue

Disk Queue (bytes)

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

This field is not applicable to agentless vSphere jobs.

Queued Replication (bytes)

The total number of replication bytes in the source queue

This field is not applicable to agentless vSphere jobs.

Sent Replication (bytes)

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

Sent Compressed Replication (bytes)

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

Sent Mirror (bytes)

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

Sent Compressed Mirror (bytes)

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

Skipped Mirror (bytes)

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

This field is not applicable to agentless vSphere jobs.

Remaining Mirror (bytes)

The total number of mirror bytes that are remaining to be sent from the source to the target

Queued Replication (ops)

The total number of replication operations in the queue

Connected Since

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

This field is not applicable to agentless vSphere jobs.

Bandwidth Limit (Kbps)

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

This field is not applicable to agentless vSphere jobs.

Specific connection items

Replication set

The name of the replication set the connection is using

Connection ID

The incremental counter used to number connections. The number is incremented when a connection is created. It is also incremented by internal actions, such as an auto-disconnect and auto-reconnect. The lowest available number (as connections are created, stopped, deleted, and so on) will always be used. The counter is reset to one each time the Double-Take service is restarted.

Transmit mode

Different jobs have different transmit mode types. The following list shows all possible transmit mode types.

- **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.
- **Restore Required**—The data on the source and target do not match because of a failover condition. Restore the data from the target back to the source. If you want to

discard the changes on the target, you can remirror to resynchronize the source and target.

- **Snapshot Reverted**—The data on the source and target do not match because a snapshot has been applied on the target. Restore the data from the target back to the source. If you want to discard the changes on the target, you can remirror to resynchronize the source and target.
- **Busy**—The source is low on memory causing a delay in getting the state of the data on the target.
- **Not Loaded**—Double-Take target functionality is not loaded on the target server. This may be caused by an activation code error.
- **Not Ready**—For a full server to ESX appliance job, the Linux drivers have not yet completed loading on the target.
- **Unknown**—The console cannot determine the status.

This field is not applicable to agentless vSphere jobs.

Target route

The target IP address and port

Compression

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

Bandwidth limit (Kbps)

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

This field is not applicable to agentless vSphere jobs.

Connected since

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

This field is not applicable to agentless vSphere jobs.

Mirror status

Different jobs have different mirror status types. The following list shows all possible mirror status types.

- **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.
- **Restoring**—Data is being restored from the target to the source.
- **Unknown**—The console cannot determine the status.

Mirror percent complete

The percentage of the mirror that has been completed

Mirror remaining (bytes)

The total number of mirror bytes that are remaining to be sent from the source to the target

Mirror skipped (bytes)

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

This field is not applicable to agentless vSphere jobs.

Queue mirror (ops)

The total number of mirror operations in the queue

Sent mirror (bytes)

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

Sent compressed mirror (bytes)

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

Replication status

Different jobs have different replication status types. The following list shows all possible replication status types.

- **Replicating**—Data is being replicated to the target.
- **In Progress**—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 (bytes)

The total number of replication bytes in the source queue

This field is not applicable to agentless vSphere jobs.

Sent replication (bytes)

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

Sent compressed replication (bytes)

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

Sent (bytes)

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

Sent compressed (bytes)

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**.

Intermediate queue (bytes)

The total amount of memory being used by the operations buffer queue

Disk queue (bytes)

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

This field is not applicable to agentless vSphere jobs.

Queued (ops)

The total number of mirror and replication operations that are in the source queue

Source Path

The location of the data on the source that is being protected

Target Path

The location on the target where the source replica data is located

Usage type

- **Normal**—The replication set type used for all job types except GeoCluster
- **GeoCluster Replicated Disk**—The replication set type for a GeoCluster job
- **Not Determined**—The replication set type could not be determined

Contains

The number of files and directories contained in the replication set

Total size

The amount of data contained in the replication set

Last calculated

The date and time the size of the replication set was last calculated

Path

The path including volume, drive, directory, file, and/or wild card

Attributes

The attributes that define the path.

- Inc—The specified path is included in the replication set
- Exc—The specified path is not included in the replication set
- Rec—The rule is automatically applied to the subdirectories of the specified path

Replication sets folder items

Name

The name of the replication set

In Use

Specifies if the replication set is being used by a connection

Last Calculated

The date and time the size of the replication set was last calculated

Size (bytes)

The amount of data contained in the replication set

Files

The number of files contained in the replication set

Directories

The number of directories contained in the replication set

Specific replication set items

Name

The name of the replication set

Usage type

- **Normal**—The replication set type used for all job types except GeoCluster
- **GeoCluster Replicated Disk**—The replication set type for a GeoCluster job
- **Not Determined**—The replication set type could not be determined

In Use

Specifies if the replication set is being used by a connection

Contains

The number of files and directories contained in the replication set

Total size

The amount of data contained in the replication set

Last calculated

The date and time the size of the replication set was last calculated

Path

The path including volume, drive, directory, file, and/or wild card

Attributes

The attributes that define the path.

- Inc—The specified path is included in the replication set
- Exc—The specified path is not included in the replication set
- Rec—The rule is automatically applied to the subdirectories of the specified path

Target connection entries folder items

Source Name

The name of the source the connection is using

Source Address

The source IP address and port the connection is using

Replication Set

The name of the replication set

Connection ID

The incremental counter used to number connections. The number is incremented when a connection is created. It is also incremented by internal actions, such as an auto-disconnect and auto-reconnect. The lowest available number (as connections are created, stopped, deleted, and so on) will always be used. The counter is reset to one each time the Double-Take service is restarted.

Target Data State

- **Disconnected**—The target connection entry is not associated with a connection. This may happen if you delete a stopped job. These disconnected target connections cannot be deleted. They will be reused when a new connection (same job type, same servers) is created.
- **OK**—The data on the target is in a good state.
- **Mirroring**—The target is in the middle of a mirror process. The data will not be in a good state until the mirror is complete.
- **Mirror Required**—The data on the target is not in a good state because a remirror is required. This may be caused by an incomplete or stopped mirror or an operation may have been dropped on the target.
- **Restore Required**—The data on the source and target do not match because of a failover condition. Restore the data from the target back to the source. If you want to discard the changes on the target, you can remirror to resynchronize the source and target.
- **Snapshot Reverted**—The data on the source and target do not match because a snapshot has been applied on the target. Restore the data from the target back to the source. If you want to discard the changes on the target, you can remirror to resynchronize the source and target.
- **Busy**—The source is low on memory causing a delay in getting the state of the data on the target.
- **Not Loaded**—Double-Take target functionality is not loaded on the target server. This may be caused by an activation code error.
- **Not Ready**—For a full server to ESX appliance job, the Linux drivers have not yet completed loading on the target.
- **Unknown**—The console cannot determine the status.

This field is not applicable to agentless vSphere jobs.

Target Status

- **OK**—The target machine is active and online.
- **Not Loaded**—The target module is not loaded on the target. (For example, the activation code is invalid.)

- **Paused**—The target machine is paused by user intervention.
- **Retrying**—The target machine is retrying operations for the connection.

Specific target connection items

Source name

The name of the source the connection is using

Source address

The source IP address and port the connection is using

Replication set

The name of the replication set

Connection ID

The incremental counter used to number connections. The number is incremented when a connection is created. It is also incremented by internal actions, such as an auto-disconnect and auto-reconnect. The lowest available number (as connections are created, stopped, deleted, and so on) will always be used. The counter is reset to one each time the Double-Take service is restarted.

Target data state

- **Disconnected**—The target connection entry is not associated with a connection. This may happen if you delete a stopped job. These disconnected target connections cannot be deleted. They will be reused when a new connection (same job type, same servers) is created.
- **OK**—The data on the target is in a good state.
- **Mirroring**—The target is in the middle of a mirror process. The data will not be in a good state until the mirror is complete.
- **Mirror Required**—The data on the target is not in a good state because a remirror is required. This may be caused by an incomplete or stopped mirror or an operation may have been dropped on the target.
- **Restore Required**—The data on the source and target do not match because of a failover condition. Restore the data from the target back to the source. If you want to discard the changes on the target, you can remirror to resynchronize the source and target.
- **Snapshot Reverted**—The data on the source and target do not match because a snapshot has been applied on the target. Restore the data from the target back to the source. If you want to discard the changes on the target, you can remirror to resynchronize the source and target.

- **Busy**—The source is low on memory causing a delay in getting the state of the data on the target.
- **Not Loaded**—Double-Take target functionality is not loaded on the target server. This may be caused by an activation code error.
- **Not Ready**—For a full server to ESX appliance job, the Linux drivers have not yet completed loading on the target.
- **Unknown**—The console cannot determine the status.

This field is not applicable to agentless vSphere jobs.

Target status

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

Schedule

Any configured snapshot schedule

Next scheduled snapshot

The date and time of the next scheduled snapshot, if any

Time Taken

The date and time of the listed snapshot

Type

- **Automatic**—This snapshot was taken automatically by Double-Take.
- **Manual**—This snapshot was taken manually by a user.
- **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.

ID

The snapshot ID

Target States

The state of the target at the time of the snapshot

Error codes

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

-1 Unknown error code (generated when a command failed but the failure is not linked to a pre-defined error code)

-101 Invalid parameter was supplied

-102 Command is not a valid or the syntax is incorrect

-103 Double-Take source module is not loaded

-104 No Double-Take source identified

-105 Double-Take target module is not loaded

-106 Connection already established

-107 Connection does not exist

-108 Mirror currently active

-109 Server does not exist or could not be located

-110 Server is not responding

-111 Double-Take is running

-112 Unknown connection error

-113 Mirror already active

-114 Date is invalid - valid format is mm/dd/yy

-115 Time is invalid - valid format is hh:mm

-116 Invalid option supplied

-117 Mirror is not paused

-118 Connection is not paused

-119 Connection does not exist

-120 Connection already connected

-121 Mirror is not running

-122 Job exists

-123 Job does not exist

-124 No job has been selected

-125 Connection is replicating

-126 Connection is not replicating

-127 Job is enabled

-128 Schedule is not defined

-129 Job is changed

-130 Job is in use

-131 No Double-Take target identified

-132 Memory is low

-133 Memory is sufficient

-134 Replication is pending

-135 Invalid option supplied

-136 Job replication rule does not exist

-137 Mirror queue is full

-138 Insufficient security access

-139 Schedule command is invalid

-140 Source path is invalid

-141 Job is not changed

-142 Insufficient source security access

-143 Invalid statistics file

-144 Job not saved

-145 Connection failed

-146 Cleaner option is not enabled

-147 Target mirror capacity high threshold is met

-148 Target mirror capacity low threshold is met

-149 New option applied

-150 Target is restarted

-151 Replication is out of memory

-152 Write access is blocked on the volume

-153 Transmission is paused

-154 Transmission is active

-155 Target does not support the command

-156 Command conversion to accommodate a different Double-Take version has failed

-157 Incompatible source and target Double-Take versions

-158 Incompatible source and target operating system versions

-159 NAS server to non-NAS server is not a supported configuration

-160 Target module is not loaded

-161 Operation or command is not supported

-162 Target is paused

-163 Target is pending

-164 Target is active

-165 Target is retrying operations

-166 Target is no longer retrying operations

-167 Restore required state is unknown

-168 Not a valid failover or cutover source

-169 Failover or cutover login failed

-170 Feature is not supported

-171 Command is not supported

-172 Target queue log file error

-173 Target disk is full

-174 Target disk has sufficient disk space

-175 Error reading from or writing to the queue log file

-176 Memory-based queue is in use

-177 Disk-based queue is in use

-178 Restore is required

-179 ID the driver supplied to the service is invalid

-180 Child path is blocked

-181 Parent path is blocked

-182 Target path blocking is disabled

-183 Connection ID specified is invalid

-184 No command objects are in the queue

-185 Target is discarding operations from the target queue

-186 Target is not discarding operations from the target queue

-187 Schedule is paused

-188 Schedule is resumed

-189 Target state has changed

-190 Target name has changed

-191 Acknowledgement queue has been updated

-201 Monitor name exists

-202 Monitor name does not exist

-203 Monitor configuration exists

-204 Monitor configuration does not exist

-205 Monitor configuration is in use

-206 Monitor configuration is not in use

-207 Source is online

-208 Source is offline

-209 Server is not failed over

-210 Server is failed over

-211 Server is not being monitored

-212 Failback is in progress

-213 IP address placeholders on the target are unavailable

-214 Target NIC was not found

-215 Source module is not loaded

-216 Failed to set the source state

-217 Unable to ping source

-218 Invalid argument

-219 Recovery is busy

-220 Invalid command

-221 Recovery is started

-222 Script failed to start

-223 Script timeout met

-224 No replication timeout met - connection is bad

-225 Invalid path

-226 Kernel module is not loaded

-227 System dump has failed

-228 Response is null

-229 Object stream is not OK

-230 Transactional NTFS (TxF) SavePoints (intermediate rollback points) are not supported

-231 Data overload

-2001 Transform initialization failed

-2002 General transform failure

-2003 Transform volume count

-2004 Transform missing source

-2005 Transform missing target

-2101 Network controller initialization failed

-2102 General network controller failure

-2103 Network controller already started

-2104 No socket on the network controller

-2105 Listen failure on the network controller

-2201 Error communicating with e-mail server

-2202 Error connecting to e-mail server

-2203 E-mail notification is disabled

-2204 E-mail notification is enabled

-2205 E-mail notification requires Internet Explorer version 5.0 and WMI (E-mail notification no longer requires Internet 5.0 or later.)

-2206 E-mail notification requires Internet Explorer version 5.0 (E-mail notification no longer requires Internet Explorer 5.0 or later.)

-2207 Error sending e-mail

-2208 Error sending test e-mail

-2209 WMI error connecting to e-mail server

-2210 E-mail notification requires WMI

-2211 Event Viewer settings for e-mail notification are invalid

-2212 E-mail notification setting is invalid

-2213 E-mail notification address exists

-2214 E-mail notification alert ID is invalid

-2215 E-mail notification format is invalid

-2216 E-mail notification address does not exist

-2217 E-mail notification address notification list is empty

-2218 E-mail warning is not set

-2219 E-mail test warning is not set

-2200 E-mail notification is functioning properly

-2301 Bandwidth limiting time exists

-2302 Bandwidth limiting name exists

-2303 Bandwidth limit not found

-2304 Bandwidth limit day is invalid

-2305 Bandwidth limit label is invalid

-2401 Snapshot module is not loaded

-2402 Error reading the snapshot .dll

-2403 Snapshot not found

-2404 No snapshot connections found

-2405 Snapshot revert completed

-2406 Snapshot revert is in progress

-2501 Full server functionality is disabled

-2502 No full server interface available

-3001 Refused target mode - Small Business Server

-3002 Refused target mode - Double-Take Move

-3003 Refused target mode - Duplicate code

-3004 Refused target mode - Double-Take Cloud

Windows Event messages

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

1. To access the Event Viewer, select **Administrative Tools**, **Event Viewer**.
2. Select the **Application** or **System** log. See your Windows reference guide or online help for details on the information provided for each event.
3. To view a detailed description, double-click an event.

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

For a complete list of Double-Take events, see *Event messages* on page 718.

Event messages

The following table identifies the Double-Take events. The event ID is followed by the event message. Below the ID and message you will find the following information.

- **Event log**—This identifies if the message will be found in the Application or System event log.
- **Source**—This identifies the Source in the event log.
- **Type or Level**—This identifies the Type (Windows 2003) or Level (Windows 2008 and 2012) in the event log.
- **Required response**—This identifies the required action, if any, you should take if you get this message.
- **SCOM alert**—This identifies if a SCOM alert rule for the message is enabled, by default, in the Double-Take Management Pack. If there is no pre-defined rule for this message, that will be indicated. See *Microsoft Systems Center Operations Manager 2007* on page 799 for details on the Management Pack.

Double-Take Availability and Double-Take Move share the same set of event messages. Some messages apply to one product, some to the other, and some to both. For messages that apply to both, the Double-Take Availability terminology is used. For example, message 5100 indicates failover completed. This same message will also be seen when cutover is completed.

Variables that are dynamically updated in a generated message are designated by a percent symbol followed by a number. For example, the message "The evaluation period expires in %1 day(s)" will have a number automatically inserted for %1, so the message you see might be "The evaluation period expires in 12 day(s)." Variables are used for things like server names, error codes, numbers, and so on.

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

Event log—Application

Source—Double-Take

Type or Level—Error

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

SCOM alert—Enabled

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

Event log—Application

Source—Double-Take

Type or Level—Information

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

SCOM alert—Enabled

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

Event log—Application

Source—Double-Take

Type or Level—Information

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

SCOM alert—Disabled

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

Event log—Application

Source—Double-Take

Type or Level—Warning

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

SCOM alert—Enabled

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

Event log—Application

Source—Double-Take

Type or Level—Error

User action required—Verify your activation code has been entered correctly.

SCOM alert—Enabled

6: Evaluation period ends today at %1.

Event log—Application

Source—Double-Take

Type or Level—Warning

User action required—Contact your vendor to purchase either a single or site license.

SCOM alert—Enabled

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

Event log—Application

Source—Double-Take

Type or Level—Error

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

SCOM alert—Enabled

100: Critical Error: %1 line %2, %3.

Event log—Application

Source—Double-Take

Type or Level—Error

User action required—Contact technical support with the details from this message.

SCOM alert—Enabled

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

Event log—Application

Source—Double-Take

Type or Level—Error

User action required—Restart the Double-Take service. Contact technical support if this event occurs repeatedly.

SCOM alert—Enabled

200: ExchFailover failover from %1 to %2 was started in commit mode. See log file %3 for details.

Event log—Application

Source—ExchFailover

Type or Level—Information

User action required—See the specific log message for additional details.

SCOM alert—Enabled

201: ExchFailover failover from %1 to %2 was started in test mode. See log file %3 for details.

Event log—Application

Source—ExchFailover

Type or Level—Information

User action required—See the specific log message for additional details.

SCOM alert—Enabled

202: ExchFailover failback to %1 from %2 was started in commit mode. See log file %3 for details.

Event log—Application

Source—ExchFailover

Type or Level—Information

User action required—See the specific log message for additional details.

SCOM alert—Enabled

203: ExchFailover failback to %1 from %2 was started in test mode. See log file %3 for details.

Event log—Application

Source—ExchFailover

Type or Level—Information

User action required—See the specific log message for additional details.

SCOM alert—Enabled

204: ExchFailover setup started for server %1. See log file %2 for details.

Event log—Application

Source—ExchFailover

Type or Level—Information

User action required—See the specific log message for additional details.

SCOM alert—Enabled

205: ExchFailover was unable to open the default log file. A new log file has been created. All messages will be log in %1.

Event log—Application

Source—ExchFailover

Type or Level—Error

User action required—See the specific log message for additional details.

SCOM alert—Enabled

210: ExchFailover completed. Moved %1 users in %2 mail stores in %3 seconds. Check log file %4 for details.

Event log—Application

Source—ExchFailover

Type or Level—Success

User action required—See the specific log message for additional details.

SCOM alert—Enabled

211: ExchFailover completed with warnings. Moved %1 users in %2 mail stores in %3 seconds. Check log file %4 for details.

Event log—Application

Source—ExchFailover

Type or Level—Warning

User action required—See the specific log message for additional details.

SCOM alert—Enabled

212: ExchFailover completed. Tested %1 users in %2 mail stores in %3 seconds. Check log file %4 for details.

Event log—Application

Source—ExchFailover

Type or Level—Success

User action required—See the specific log message for additional details.

SCOM alert—Enabled

213: ExchFailover completed with warnings. Moved %1 users in %2 mail stores in %3 seconds. Check log file %4 for details.

Event log—Application

Source—ExchFailover

Type or Level—Warning

User action required—See the specific log message for additional details.

SCOM alert—Enabled

214: ExchFailover setup completed. Updated %1 mail stores in %2 seconds. Check log file %3 for details.

Event log—Application

Source—ExchFailover

Type or Level—Success

User action required—See the specific log message for additional details.

SCOM alert—Enabled

220: ExchFailover start failed. Could not open log file: %1.

Event log—Application

Source—ExchFailover

Type or Level—Error

User action required—Restart failover. Contact technical support if this event occurs again.

SCOM alert—Enabled

221: ExchFailover start failed. Invalid command line arguments. See log file %1 for details.

Event log—Application

Source—ExchFailover

Type or Level—Error

User action required—See the specific log message for additional details.

SCOM alert—Enabled

222: ExchFailover start failed. Double-Take is not licensed on this machine.

Event log—Application

Source—ExchFailover

Type or Level—Error

User action required—Verify your activation code has been entered correctly and contact technical support.

SCOM alert—Enabled

223: ExchFailover start failed due to an Active Directory error.

Event log—Application

Source—ExchFailover

Type or Level—Error

User action required—Restart failover. Contact technical support if this event occurs again.

SCOM alert—Enabled

224: ExchFailover failed to find one (or both) of the Exchange servers. Check the server names. This can also occur if the process does not have sufficient privileges to access Active Directory.

Event log—Application

Source—ExchFailover

Type or Level—Error

User action required—Verify the Exchange server names and the account has sufficient privileges to update Active Directory.

SCOM alert—Enabled

1000: An exception occurred: %1

Event log—Application

Source—DTCounters

Type or Level—Error

User action required—Run the installation and select Repair. Contact technical support if this event occurs again.

SCOM alert—Enabled

1001: The Double-Take counter DLL could not initialize the statistics handler object to gather performance data.

Event log—Application

Source—DTCounters

Type or Level—Error

User action required—Run the installation and select Repair. Contact technical support if this event occurs again.

SCOM alert—Enabled

1002: The Double-Take counter DLL could not map shared memory file containing the performance data.

Event log—Application

Source—DTCounters

Type or Level—Error

User action required—Run the installation and select Repair. Contact technical support if this event occurs again.

SCOM alert—Enabled

1003: The Double-Take counter DLL could not open the "Performance" key in the Double-Take section of the registry.

Event log—Application

Source—DTCounters

Type or Level—Error

User action required—Run the installation and select Repair. Contact technical support if this event occurs again.

SCOM alert—Enabled

1004: The Double-Take counter DLL could not read the "First Counter" value under the Double-Take\Performance Key.

Event log—Application

Source—DTCounters

Type or Level—Error

User action required—Run the installation and select Repair. Contact technical support if this event occurs again.

SCOM alert—Enabled

1005: The Double-Take counter DLL read the "First Help" value under the Double-Take\Performance Key.

Event log—Application

Source—DTCounters

Type or Level—Error

User action required—Run the installation and select Repair. Contact technical support if this event occurs again.

SCOM alert—Enabled

1006: The Double-Take counter DLL could not create event handler for the worker thread.

Event log—Application

Source—DTCounters

Type or Level—Error

User action required—Run the installation and select Repair. Contact technical support if this event occurs again.

SCOM alert—Enabled

4000: Kernel was successfully started.

Event log—Application

Source—Double-Take

Type or Level—Information

User action required—No action required.

SCOM alert—Enabled

4001: Target service was successfully started.

Event log—Application

Source—Double-Take

Type or Level—Information

User action required—No action required.

SCOM alert—Enabled

4002: Source service was successfully started.

Event log—Application

Source—Double-Take

Type or Level—Information

User action required—No action required.

SCOM alert—Enabled

4003: Source service was successfully stopped.

Event log—Application

Source—Double-Take

Type or Level—Information

User action required—No action required.

SCOM alert—Enabled

4004: Target service was successfully stopped.

Event log—Application

Source—Double-Take

Type or Level—Information

User action required—No action required.

SCOM alert—Enabled

4005: Kernel was successfully stopped.

Event log—Application

Source—Double-Take

Type or Level—Information

User action required—No action required.

SCOM alert—Enabled

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

Event log—Application

Source—Double-Take

Type or Level—Warning

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

SCOM alert—Enabled

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

Event log—Application

Source—Double-Take

Type or Level—Information

User action required—No action required.

SCOM alert—Enabled

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

Event log—Application

Source—Double-Take

Type or Level—Information

User action required—No action required.

SCOM alert—Enabled

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

Event log—Application

Source—Double-Take

Type or Level—Information

User action required—No action required.

SCOM alert—Disabled

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

Event log—Application

Source—Double-Take

Type or Level—Error

User action required—Manually reestablish the job to target connection.

SCOM alert—Enabled

4014: Service has started network transmission.

Event log—Application

Source—Double-Take

Type or Level—Information

User action required—No action required.

SCOM alert—Disabled

4015: Service has stopped network transmission.

Event log—Application

Source—Double-Take

Type or Level—Information

User action required—No action required.

SCOM alert—Enabled

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

Event log—Application

Source—Double-Take

Type or Level—Information

User action required—No action required.

SCOM alert—Disabled

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

Event log—Application

Source—Double-Take

Type or Level—Information

User action required—No action required.

SCOM alert—Enabled

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

Event log—Application

Source—Double-Take

Type or Level—Warning

User action required—Perform a restoration.

SCOM alert—Enabled

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

Event log—Application

Source—Double-Take

Type or Level—Information

User action required—No action required.

SCOM alert—Enabled

4020: Service has paused a mirror to %1 (%2) for Replication Set %3, ID: %4

Event log—Application

Source—Double-Take

Type or Level—Information

User action required—No action required.

SCOM alert—Enabled

4021: Service has resumed a mirror to %1 (%2) for Replication Set %3, ID: %4

Event log—Application

Source—Double-Take

Type or Level—Information

User action required—No action required.

SCOM alert—Disabled

4022: Service has stopped a mirror to %1 for Replication Set %2, ID: %3, %4

Event log—Application

Source—Double-Take

Type or Level—Information

User action required—No action required.

SCOM alert—Enabled

4023: Service has completed a mirror to %1 %2 for Replication Set %3, ID: %4, %5

Event log—Application

Source—Double-Take

Type or Level—Information

User action required—No action required.

SCOM alert—Enabled

4024: Service has started Replication to %1 (%2) for Replication Set %3, ID: %4

Event log—Application

Source—Double-Take

Type or Level—Information

User action required—No action required.

SCOM alert—Disabled

4025: Service has stopped Replication to %1 (%2) for Replication Set %3, ID: %4

Event log—Application

Source—Double-Take

Type or Level—Information

User action required—No action required.

SCOM alert—Enabled

4026: The target has been paused due to user intervention.

Event log—Application

Source—Double-Take

Type or Level—Information

User action required—No action required.

SCOM alert—Enabled

4027: The target has been resumed due to user intervention.

Event log—Application

Source—Double-Take

Type or Level—Information

User action required—No action required.

SCOM alert—Enabled

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

Event log—Application

Source—Double-Take

Type or Level—Warning

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

SCOM alert—Enabled

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

Event log—Application

Source—Double-Take

Type or Level—Warning

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

SCOM alert—Enabled

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

Event log—Application

Source—Double-Take

Type or Level—Error

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

SCOM alert—Enabled

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

Event log—Application

Source—Double-Take

Type or Level—Error

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

SCOM alert—Enabled

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

Event log—Application

Source—Double-Take

Type or Level—Error

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

SCOM alert—Enabled

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

Event log—Application

Source—Double-Take

Type or Level—Error

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

SCOM alert—Enabled

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

Event log—Application

Source—Double-Take

Type or Level—Error

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

SCOM alert—Enabled

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

Event log—Application

Source—Double-Take

Type or Level—Error

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

SCOM alert—Enabled

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

Event log—Application

Source—Double-Take

Type or Level—Error

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

SCOM alert—Enabled

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

Event log—Application

Source—Double-Take

Type or Level—Error

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

SCOM alert—Enabled

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

Event log—Application

Source—Double-Take

Type or Level—Error

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

SCOM alert—Enabled

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

Event log—Application

Source—Double-Take

Type or Level—Error

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

SCOM alert—Enabled

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

Event log—Application

Source—Double-Take

Type or Level—Error

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

SCOM alert—Enabled

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

Event log—Application

Source—Double-Take

Type or Level—Error

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

SCOM alert—Enabled

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

Event log—Application

Source—Double-Take

Type or Level—Error

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

SCOM alert—Enabled

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

Event log—Application

Source—Double-Take

Type or Level—Error

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

SCOM alert—Enabled

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

Event log—Application

Source—Double-Take

Type or Level—Error

User action required—Contact technical support if this error persists.

SCOM alert—Enabled

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

Event log—Application

Source—Double-Take

Type or Level—Warning

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

SCOM alert—Enabled

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

Event log—Application

Source—Double-Take

Type or Level—Error

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

SCOM alert—Enabled

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

Event log—Application

Source—Double-Take

Type or Level—Information

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

SCOM alert—Enabled

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

Event log—Application

Source—Double-Take

Type or Level—Information

User action required—No action required.

SCOM alert—Disabled

4049: Service has paused a delete orphans task to %1 (%2) for Replication Set %3, ID: %4

Event log—Application

Source—Double-Take

Type or Level—Information

User action required—No action required.

SCOM alert—Disabled

4050: Service has resumed a delete orphans task to %1 (%2) for Replication Set %3, ID: %4

Event log—Application

Source—Double-Take

Type or Level—Information

User action required—No action required.

SCOM alert—Disabled

4051: Service has stopped a delete orphans task to %1 (%2) for Replication Set %3, ID: %4

Event log—Application

Source—Double-Take

Type or Level—Information

User action required—No action required.

SCOM alert—Disabled

4052: Service has completed a delete orphans task to %1 (%2) for Replication Set %3, ID: %4

Event log—Application

Source—Double-Take

Type or Level—Information

User action required—No action required.

SCOM alert—Disabled

4053: Service has started a restore task to %1 (%2) for Replication Set %3, ID: %4

Event log—Application

Source—Double-Take

Type or Level—Information

User action required—No action required.

SCOM alert—Disabled

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

Event log—Application

Source—Double-Take

Type or Level—Information

User action required—No action required.

SCOM alert—Disabled

4055: Service has resumed a restore task to %1 (%2) for Replication Set %3, ID: %4

Event log—Application

Source—Double-Take

Type or Level—Information

User action required—No action required.

SCOM alert—Disabled

4056: Service has stopped a restore task to %1 (%2) for Replication Set %3, ID: %4

Event log—Application

Source—Double-Take

Type or Level—Information

User action required—No action required.

SCOM alert—Enabled

4057: Service has completed a restore task to %1 (%2) for Replication Set %3, ID: %4

Event log—Application

Source—Double-Take

Type or Level—Information

User action required—No action required.

SCOM alert—Enabled

4058: Service has started a verification task to %1 (%2) for Replication Set %3, ID: %4

Event log—Application

Source—Double-Take

Type or Level—Information

User action required—No action required.

SCOM alert—Disabled

4059: Service has paused a verification task to %1 (%2) for Replication Set %3, ID: %4

Event log—Application

Source—Double-Take

Type or Level—Information

User action required—No action required.

SCOM alert—Disabled

4060: Service has resumed a verification task to %1 (%2) for Replication Set %3, ID: %4

Event log—Application

Source—Double-Take

Type or Level—Information

User action required—No action required.

SCOM alert—Disabled

4061: Service has stopped a verification task to %1 (%2) for Replication Set %3, ID: %4

Event log—Application

Source—Double-Take

Type or Level—Information

User action required—No action required.

SCOM alert—Enabled

4062: Service has completed a verification task to %1 (%2) for Replication Set %3, ID: %4

Event log—Application

Source—Double-Take

Type or Level—Information

User action required—No action required.

SCOM alert—Enabled

4063: Bandwidth limit to %1 (%2) has changed to %3.

Event log—Application

Source—Double-Take

Type or Level—Information

User action required—No action required.

SCOM alert—Enabled

4064: Bandwidth limit to %1 (%2) is now in the "%3" period at %4.

Event log—Application

Source—Double-Take

Type or Level—Information

User action required—No action required.

SCOM alert—Disabled

4065: Target data state for connection %1 from %2 (%3) has changed because %4.

Event log—Application

Source—Double-Take

Type or Level—Warning

User action required—No action required.

SCOM alert—Enabled

4066: The product code requires a virtual server environment. The product functionality has been disabled.

Event log—Application

Source—Double-Take

Type or Level—Error

User action required—The activation code you are using is for the Virtual SystemsTM edition. This code will not work on non-virtual server environments.

SCOM alert—Enabled

4067: No replication ops have been received from the driver for an extended period of time.

Event log—Application

Source—Double-Take

Type or Level—Error

User action required—Check other messages for errors with the Double-Take drivers, and correct as required. If there are no driver messages, verify that your drives are connected to the source. If this error persists, contact technical support.

SCOM alert—Enabled

4068: Failed to write to a replicating volume.

Event log—Application

Source—Double-Take

Type or Level—Error

User action required—Reboot the source server. Contact technical support if this event occurs again.

SCOM alert—Enabled

4069: The option MoveOrphansDir has been updated because it was missing or empty.

Event log—Application

Source—Double-Take

Type or Level—Warning

User action required—No action required.

SCOM alert—Enabled

4070: An error occurred while reading data for connection %1. All data needs to be remirrored. See the log for details.

Event log—Application

Source—Double-Take

Type or Level—Error

User action required—Initiate a remirror to guarantee data integrity. Contact technical support if this event occurs repeatedly.

SCOM alert—Enabled

4071: Received network message with invalid checksum.

Event log—Application

Source—Double-Take

Type or Level—Warning

User action required—Initiate a remirror to guarantee data integrity. Contact technical support if this event occurs repeatedly.

SCOM alert—Enabled

4072: QueueSizeAlertThreshold of %1% has been exceeded.

Event log—Application

Source—Double-Take

Type or Level—Warning

User action required—If the queue reaches capacity, Double-Take will automatically begin the auto-disconnect process. If you see this message repeatedly, you may want to consider a larger queue or upgrading your server hardware to keep up with the amount of data changes in your environment.

SCOM alert—Enabled

4096: The registry parameter %2 is unknown.

Event log—System

Source—RepDrv

Type or Level—Warning

User action required—Delete the parameter and report this issue to technical support.

SCOM alert—Enabled

4097: Failed to initialize WMI support. The last Word in the Data Window is the NT status code.

Event log—System

Source—RepDrv, RepKap, RepHsm, or RepSis

Type or Level—Warning

User action required—No action required.

SCOM alert—Enabled

4097: The file system filter failed to load. Replication will not occur. Reboot your server and contact technical support if this error occurs again. The last Word in the Data window is the NT status code.

Event log—System

Source—RepDrv

Type or Level—Error

User action required—Reboot your server and contact technical support if this event occurs again.

SCOM alert—Enabled

4098: The registry parameters failed to load, so the default configuration values will be used. The last Word in the Data window is the NT status code.

Event log—System

Source—RepKap

Type or Level—Warning

User action required—No action required.

SCOM alert—Enabled

4098: The control device %2 was not created. Communication with the service will be disabled. Reboot the server and contact technical support if this error occurs again. The last Word in the Data window is the NT status code.

Event log—System

Source—RepDrv, RepDac, RepKap, or RepHsm

Type or Level—Error

User action required—Reboot your server and contact technical support if this event occurs again.

SCOM alert—Enabled

4099: The driver detected a hard link for a file on drive %2. Hard links are not supported. Changes to this file will not be replicated.

Event log—System

Source—RepDrv

Type or Level—Warning

User action required—Hard links are not supported.

SCOM alert—Enabled

4099: The driver failed to register with filter manager. Reboot the server and contact technical support if this error occurs again. The last Word in the Data window is the NT status code.

Event log—System

Source—RepKap

Type or Level—Error

User action required—Reboot your server and contact technical support if this event occurs again.

SCOM alert—Enabled

4100: The versions of the driver and the filter driver do not match. Replication will not occur. Reboot your server. If this error occurs again, reinstall the software. Contact technical support if this error occurs after the software has been reinstalled. The last three Words in the Data window are the NT status code and the driver version numbers.

Event log—System

Source—RepDrv

Type or Level—Error

User action required—Reboot your server. Reinstall the software if this event occurs again. Contact technical support if this event occurs after reinstalling the software.

SCOM alert—Enabled

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

Event log—Application

Source—Double-Take

Type or Level—Warning

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

SCOM alert—Enabled

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

Event log—Application

Source—Double-Take

Type or Level—Warning

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

SCOM alert—Enabled

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

Event log—Application

Source—Double-Take

Type or Level—Warning

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

SCOM alert—Enabled

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

Event log—Application

Source—Double-Take

Type or Level—Warning

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

SCOM alert—Enabled

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

Event log—Application

Source—Double-Take

Type or Level—Information

User action required—No action required.

SCOM alert—Disabled

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

Event log—Application

Source—Double-Take

Type or Level—Error

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

SCOM alert—Enabled

4155: The service was unable to complete a file system operation in the allotted time. See the log files for error conditions. After correcting the problem, remirror or perform a verification with remirror to synchronize the changes.

Event log—Application

Source—Double-Take

Type or Level—Warning

User action required—Correct the file system error and then remirror or perform a verification with remirror to synchronize the changes.

SCOM alert—Enabled

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

Event log—Application

Source—Double-Take

Type or Level—Information

User action required—No action required.

SCOM alert—Disabled

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

Event log—Application

Source—Double-Take

Type or Level—Warning

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

SCOM alert—Enabled

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

Event log—Application

Source—Double-Take

Type or Level—Information

User action required—No action required.

SCOM alert—Disabled

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

Event log—Application

Source—Double-Take

Type or Level—Information

User action required—No action required.

SCOM alert—Disabled

4204: Error running in band script: %1

Event log—Application

Source—Double-Take

Type or Level—Error

User action required—Review the task and its associated script(s) for syntax errors.

SCOM alert—Enabled

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

Event log—Application

Source—Double-Take

Type or Level—Warning

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

SCOM alert—Enabled

4206: Run timeout disabled for in band script: %1

Event log—Application

Source—Double-Take

Type or Level—Warning

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

SCOM alert—Enabled

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

Event log—Application

Source—Double-Take

Type or Level—Warning

User action required—Enable task command processing.

SCOM alert—Enabled

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

Event log—Application

Source—Double-Take

Type or Level—Error

User action required—You may need to stop and restart your job.

SCOM alert—Enabled

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

Event log—Application

Source—Double-Take

Type or Level—Error

User action required—If you need to block your target paths, contact technical support.

SCOM alert—Enabled

4302: Target Path %1 has been successfully blocked

Event log—Application

Source—Double-Take

Type or Level—Information

User action required—No action required.

SCOM alert—Disabled

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

Event log—Application

Source—Double-Take

Type or Level—Warning

User action required—If you need to block your target paths, contact technical support.

SCOM alert—Enabled

4304: Target Path %1 has been successfully unblocked

Event log—Application

Source—Double-Take

Type or Level—Information

User action required—No action required.

SCOM alert—Disabled

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

Event log—Application

Source—Double-Take

Type or Level—Warning

User action required—If you need to unblock your target paths, contact technical support.

SCOM alert—Enabled

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

Event log—Application

Source—Double-Take

Type or Level—Warning

User action required—No action required.

SCOM alert—Enabled

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

Event log—Application

Source—Double-Take

Type or Level—Warning

User action required—No action required.

SCOM alert—Disabled

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

Event log—Application

Source—Double-Take

Type or Level—Error

User action required—If you need to block your target paths, contact technical support.

SCOM alert—Enabled

4400: Failed to create snapshot set for source %1 (%2) Connection ID: %3. Error: %4

Event log—Application

Source—Double-Take

Type or Level—Error

User action required—The snapshot could not be created. This may be due to a lack of disk space or memory or another reason. The error code is the Microsoft VSS error. Check your VSS documentation or contact technical support.

SCOM alert—Enabled

4401: Failed to delete automatic snapshot set for source %1 (%2) Connection ID: %3. Error: %4

Event log—Application

Source—Double-Take

Type or Level—Error

User action required—The automatic snapshot could not be deleted. This may be due to a lack of memory, the file does not exist, or another reason. The error code is the Microsoft Volume Shadow Copy error. Check your Volume Shadow Copy documentation or contact technical support.

SCOM alert—Enabled

4402: Failed to delete snapshot set for source %1 (%2) Connection ID: %3. Error: %4

Event log—Application

Source—Double-Take

Type or Level—Error

User action required—The snapshot could not be deleted. This may be due to a lack of memory, the file does not exist, or another reason. The error code is the Microsoft Volume Shadow Copy error. Check your Volume Shadow Copy documentation or contact technical support.

SCOM alert—Enabled

4403: A scheduled snapshot could not be created for source %1 (%2) Connection ID: %3. because the target data was in a bad state. A snapshot will automatically be created when the target data reaches a good state.

Event log—Application

Source—Double-Take

Type or Level—Error

User action required—No action required. A snapshot will automatically be created when the target data reaches a good state.

SCOM alert—Enabled

4404: Set snapshot schedule for source %1 (%2) connection %3 to every %4 minutes. Next snapshot: %5.

Event log—Application

Source—Double-Take

Type or Level—Information

User action required—No action required.

SCOM alert—Enabled

4405: Removed snapshot schedule for source %1 (%2) connection %3.

Event log—Application

Source—Double-Take

Type or Level—Information

User action required—No action required.

SCOM alert—Enabled

4406: Enabled snapshot schedule for source %1 (%2) connection %3.

Event log—Application

Source—Double-Take

Type or Level—Information

User action required—No action required.

SCOM alert—Enabled

4407: Disabled snapshot schedule for source %1 (%2) connection %3.

Event log—Application

Source—Double-Take

Type or Level—Information

User action required—No action required.

SCOM alert—Enabled

4408: %1 was unable to move some orphans for source %2 on connection ID %3. Check the %1 logs for further details.

Event log—Application

Source—Double-Take

Type or Level—Warning

User action required—Orphan files could not be moved. For example, the location could be out of disk space. Check the Double-Take log for more information.

SCOM alert—Enabled

4409: %3 was unable to delete some orphans for source %1 on connection ID %2. Check the %3 logs for further details.

Event log—Application

Source—Double-Take

Type or Level—Warning

User action required—Orphan files could not be deleted. Check the Double-Take log for more information.

SCOM alert—Enabled

4410: The registry hive dump failed with an of error: %1.

Event log—Application

Source—Double-Take

Type or Level—Error

User action required—Contact technical support.

SCOM alert—Enabled

4411: The Service has detected that port %1 is being %2 by the Windows Firewall.

Event log—Application

Source—Double-Take

Type or Level—Warning

User action required—The firewall port needs to be unblocked or restrictions against Double-Take removed so that Double-Take data can be transmitted.

SCOM alert—Enabled

5100: Failover completed for %1.

Event log—Application

Source—Double-Take

Type or Level—Information

User action required—No action required.

SCOM alert—Enabled

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

Event log—Application

Source—Double-Take

Type or Level—Information

User action required—No action required.

SCOM alert—Disabled

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

Event log—Application

Source—Double-Take

Type or Level—Warning

User action required—Check your source machine and initiate failover, if user intervention for failover is configured. If you bring your source machine back online without initiating failover, the failover condition met state will be canceled.

SCOM alert—Enabled

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

Event log—Application

Source—Double-Take

Type or Level—Information

User action required—No action required.

SCOM alert—Disabled

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

Event log—Application

Source—Double-Take

Type or Level—Information

User action required—No action required.

SCOM alert—Disabled

5105: Attempting to run the %1 script.

Event log—Application

Source—Double-Take

Type or Level—Information

User action required—No action required.

SCOM alert—Disabled

5106: The %1 script ran successfully.

Event log—Application

Source—Double-Take

Type or Level—Information

User action required—No action required.

SCOM alert—Disabled

5107: Error occurred in running %1 script.

Event log—Application

Source—Double-Take

Type or Level—Error

User action required—Verify that the script identified exists with the proper permissions.

SCOM alert—Enabled

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

Event log—Application

Source—Double-Take

Type or Level—Error

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

SCOM alert—Enabled

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

Event log—Application

Source—Double-Take

Type or Level—Error

User action required—The failover target did not receive an answer to its ping of the source machine. Eventually, a failover will result. Investigate possible errors (down server, network error, and so on).

SCOM alert—Enabled

5110: The %1 script "%2" is still running.

Event log—Application

Source—Double-Take

Type or Level—Information

User action required—No action required.

SCOM alert—Enabled

5200: Failback completed for %1.

Event log—Application

Source—Double-Take

Type or Level—Information

User action required—No action required.

SCOM alert—Enabled

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

Event log—Application

Source—Double-Take

Type or Level—Information

User action required—No action required.

SCOM alert—Enabled

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

Event log—Application

Source—Double-Take

Type or Level—Error

User action required—Contact technical support.

SCOM alert—Enabled

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

Event log—Application

Source—Double-Take

Type or Level—Information

User action required—No action required.

SCOM alert—Disabled

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

Event log—Application

Source—Double-Take

Type or Level—Information

User action required—No action required.

SCOM alert—Disabled

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

Event log—Application

Source—Double-Take

Type or Level—Information

User action required—No action required.

SCOM alert—Disabled

5303: The application monitor script has started successfully.

Event log—Application

Source—Double-Take

Type or Level—Information

User action required—No action required.

SCOM alert—Enabled

5304: The application monitor script has finished successfully.

Event log—Application

Source—Double-Take

Type or Level—Information

User action required—No action required.

SCOM alert—Enabled

5305: The application monitor has found the %1 service stopped.

Event log—Application

Source—Double-Take

Type or Level—Warning

User action required—Application Manager will attempt to restart the service.

SCOM alert—Enabled

5306: The application monitor has restarted the %1 service.

Event log—Application

Source—Double-Take

Type or Level—Warning

User action required—No action required.

SCOM alert—Enabled

5307: The application monitor cannot contact the server %1.

Event log—Application

Source—Double-Take

Type or Level—Error

User action required—Verify the server is running. Verify available network communications with the server.

SCOM alert—Enabled

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

Event log—Application

Source—Double-Take

Type or Level—Information

User action required—No action required.

SCOM alert—Disabled

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

Event log—Application

Source—Double-Take

Type or Level—Warning

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

SCOM alert—Enabled

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

Event log—Application

Source—Double-Take

Type or Level—Warning

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

SCOM alert—Enabled

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

Event log—Application

Source—Double-Take

Type or Level—Warning

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

SCOM alert—Enabled

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

Event log—Application

Source—Double-Take

Type or Level—Warning

User action required—If you are using Double-Take 4.4.2.1 or earlier and Windows NT 4.0, e-mail notification requires Windows Management Instrumentation (WMI) to be installed. Verify that you have it installed on the Double-Take server.

SCOM alert—Enabled

5504: Could not load LocalRS.dll (for e-mail notification).

Event log—Application

Source—Double-Take

Type or Level—Warning

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

SCOM alert—Enabled

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

Event log—Application

Source—Double-Take

Type or Level—Warning

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

SCOM alert—Enabled

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

Event log—Application

Source—Double-Take

Type or Level—Warning

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

SCOM alert—Enabled

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

Event log—Application

Source—Double-Take

Type or Level—Warning

User action required—If you are using Double-Take 4.4.2.1 or earlier and Windows NT 4.0, e-mail notification requires Windows Management Instrumentation (WMI) to be installed. Verify that you have it installed on the Double-Take server.

SCOM alert—Enabled

5508: An error occurred connecting to the WMI namespace. Check to make sure the Windows Management Instrumentation service is not disabled (error code %1).

Event log—Application

Source—Double-Take

Type or Level—Warning

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

SCOM alert—Enabled

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

Event log—Application

Source—Double-Take

Type or Level—Warning

User action required—Verify that the include categories and exclude ID list are identified and formatted correctly.

SCOM alert—Enabled

6000: %1

Event log—Application

Source—Double-Take Management Service

Type or Level—Information

User action required—This is a placeholder message for many other messages. See the specific log message for additional details.

SCOM alert—Disabled

6001: %1

Event log—Application

Source—Double-Take Management Service

Type or Level—Warning

User action required—This is a placeholder message for many other messages. See the specific log message for additional details.

SCOM alert—Enabled

6002: %1

Event log—Application

Source—Double-Take Management Service

Type or Level—Error

User action required—This is a placeholder message for many other messages. See the specific log message for additional details.

SCOM alert—Enabled

6003: A %1 job has been created. The name is "%2" and the ID is %3.

Event log—Application

Source—Double-Take Management Service

Type or Level—Information

User action required—No action required.

SCOM alert—Enabled

6004: The %1 job "%2" (ID %3) has been started.

Event log—Application

Source—Double-Take Management Service

Type or Level—Information

User action required—No action required.

SCOM alert—Enabled

6005: The %1 job "%2" (ID %3) has been stopped.

Event log—Application

Source—Double-Take Management Service

Type or Level—Information

User action required—No action required. If desired, you can restart your job.

SCOM alert—Enabled

6006: The %1 job "%2" (ID %3) has been deleted.

Event log—Application

Source—Double-Take Management Service

Type or Level—Information

User action required—No action required. If desired, you can re-create your job.

SCOM alert—Enabled

6007: The %1 operation has failed for the %2 job "%3" (ID %4).

Event log—Application

Source—Double-Take Management Service

Type or Level—Error

User action required—No action required. If desired, you can re-create your job.

SCOM alert—Enabled

6008: The %1 operation has completed successfully for the %2 job "%3" (ID %4).

Event log—Application

Source—Double-Take Management Service

Type or Level—Information

User action required—No action required.

SCOM alert—Enabled

6009: Could not log the following message:%n%1%n---%nError:%n%2.

Event log—Application

Source—Double-Take Management Service

Type or Level—Error

User action required—There is a problem with logging. Contact technical support if this event occurs again.

SCOM alert—Enabled

6010: A failover condition has been met for the %1 job "%2" (ID %3).

Event log—Application

Source—Double-Take Management Service

Type or Level—Warning

User action required—Check your source machine and initiate failover, if user intervention for failover is configured. If you bring your source machine back online without initiating failover, the failover condition met state will be canceled.

SCOM alert—Enabled

6011: The source machine (IP %1) is not responding to a ping from monitor %2 (ID %3).

Event log—Application

Source—Double-Take Management Service

Type or Level—Error

User action required—Check your source machine and initiate failover, if user intervention for failover is configured. If you bring your source machine back online without initiating failover, the source machine should start responding to the ping.

SCOM alert—Enabled

6012: The target machine (IP %1) failed to reboot.

Event log—Application

Source—Double-Take Management Service

Type or Level—Error

User action required—Reboot the target server to complete full server failover.

SCOM alert—Enabled

6050: The service has detected that port %1 is RESTRICTED in the Windows Firewall. This port is critical to the operation of the Double-Take Management Service.

Event log—Application

Source—Double-Take Management Service

Type or Level—Error

User action required—Verify the specified firewall port is open for Double-Take traffic.

SCOM alert—Enabled

6051: The service has detected that port %1 is BLOCKED in the Windows Firewall. This port is critical to the operation of the Double-Take Management Service.

Event log—Application

Source—Double-Take Management Service

Type or Level—Error

User action required—Verify the specified firewall port is open for Double-Take traffic.

SCOM alert—Enabled

6100: The job "%1" (ID %2) has started provisioning a replica for %3.

Event log—Application

Source—Double-Take Management Service

Type or Level—Information

User action required—No action required.

SCOM alert—Enabled

6101: The job "%1" (ID %2) has successfully completed provisioning a replica for %3.

Event log—Application

Source—Double-Take Management Service

Type or Level—Information

User action required—No action required.

SCOM alert—Enabled

6102: The job "%1" (ID %2) has failed to provision a replica for %3.%n%n

Event log—Application

Source—Double-Take Management Service

Type or Level—Error

User action required—Review the additional error information to identify the problem. Correct the problem and retry the operation. Contact technical support if this event occurs again.

SCOM alert—Enabled

6110: The job "%1" (ID %2) has started a %3 failover of the replica of %4.

Event log—Application

Source—Double-Take Management Service

Type or Level—Information

User action required—No action required.

SCOM alert—Enabled

6111: The job "%1" (ID %2) has successfully completed a %3 failover of the replica of %4.

Event log—Application

Source—Double-Take Management Service

Type or Level—Information

User action required—No action required.

SCOM alert—Enabled

6112: The job "%1" (ID %2) has encountered an error while performing a %3 failover of the replica of %4.%n%n

Event log—Application

Source—Double-Take Management Service

Type or Level—Error

User action required—Review the additional error information to identify the problem. Correct the problem and retry the operation. Contact technical support if this event occurs again.

SCOM alert—Enabled

6120: The job "%1" (ID %2) has started undoing the failover for the replica of %3.

Event log—Application

Source—Double-Take Management Service

Type or Level—Information

User action required—No action required.

SCOM alert—Enabled

6121: The job "%1" (ID %2) has successfully reattached the replica and resumed protecting %3.

Event log—Application

Source—Double-Take Management Service

Type or Level—Information

User action required—No action required.

SCOM alert—Enabled

6122: The job "%1" (ID %2) has encountered an error undoing the failing over for the replica of %3.%n%n

Event log—Application

Source—Double-Take Management Service

Type or Level—Error

User action required—Review the additional error information to identify the problem. Correct the problem and retry the operation. Contact technical support if this event occurs again.

SCOM alert—Enabled

6130: The job "%1" (ID %2) has started reversing the direction of the protection of %3.

Event log—Application

Source—Double-Take Management Service

Type or Level—Information

User action required—No action required.

SCOM alert—Enabled

6131: The job "%1" (ID %2) has successfully reversed the direction of the protection of %3.

Event log—Application

Source—Double-Take Management Service

Type or Level—Information

User action required—No action required.

SCOM alert—Enabled

6132: The job "%1" (ID %2) has encountered an error reversing the direction of the protection of %3.%n%n

Event log—Application

Source—Double-Take Management Service

Type or Level—Error

User action required—Review the additional error information to identify the problem. Correct the problem and retry the operation. Contact technical support if this event occurs again.

SCOM alert—Enabled

6140: The job "%1" (ID %2) is being deleted.

Event log—Application

Source—Double-Take Management Service

Type or Level—Information

User action required—No action required.

SCOM alert—Enabled

6141: The job "%1" (ID %2) has successfully been deleted.

Event log—Application

Source—Double-Take Management Service

Type or Level—Information

User action required—No action required.

SCOM alert—Enabled

6142: The job "%1" (ID %2) has encountered an error while being deleted.%n%n

Event log—Application

Source—Double-Take Management Service

Type or Level—Error

User action required—Review the additional error information to identify the problem. Correct the problem and retry the operation. Contact technical support if this event occurs again.

SCOM alert—Enabled

6150: The job "%1" (ID %2) protecting %3 has completed its mirror.

Event log—Application

Source—Double-Take Management Service

Type or Level—Information

User action required—No action required.

SCOM alert—Enabled

6210: The job "%1" (ID %2) has started a %3 failover of the replica of %4.

Event log—Application

Source—Double-Take Management Service

Type or Level—Information

User action required—No action required.

SCOM alert—Enabled

6211: The job "%1" (ID %2) has successfully completed a %3 failover of the replica of %4.

Event log—Application

Source—Double-Take Management Service

Type or Level—Information

User action required—No action required.

SCOM alert—Enabled

6212: The job "%1" (ID %2) has encountered an error while performing a %3 failover of the replica of %4.%n%n

Event log—Application

Source—Double-Take Management Service

Type or Level—Error

User action required—Review the additional error information to identify the problem. Correct the problem and retry the operation. Contact technical support if this event occurs again.

SCOM alert—Enabled

6213: A failover condition has been met for the host level job "%1" (ID %2).

Event log—Application

Source—Double-Take Management Service

Type or Level—Warning

User action required—Check your source machine and initiate failover, if user intervention for failover is configured. If you bring your source machine back online without initiating failover, the failover condition met state will be canceled.

SCOM alert—Enabled

6214: Failover monitors removed for the host level job "%1" (ID %2).

Event log—Application

Source—Double-Take Management Service

Type or Level—Information

User action required—No action required. The failover pending state has been canceled because the job has been stopped, deleted, or failed over.

SCOM alert—Enabled

6215: The job "%1" (ID %2) is protecting.

Event log—Application

Source—Double-Take Management Service

Type or Level—Information

User action required—No action required.

SCOM alert—Enabled

6220: The job "%1" (ID %2) has started undoing the failover for the replica of %3.

Event log—Application

Source—Double-Take Management Service

Type or Level—Information

User action required—No action required.

SCOM alert—Enabled

6221: The job "%1" (ID %2) has successfully reattached the replica and resumed protecting %3.

Event log—Application

Source—Double-Take Management Service

Type or Level—Information

User action required—No action required.

SCOM alert—Enabled

6222: The job "%1" (ID %2) has encountered an error undoing the failing over for the replica of %3.%n%n

Event log—Application

Source—Double-Take Management Service

Type or Level—Error

User action required—Review the additional error information to identify the problem. Correct the problem and retry the operation. Contact technical support if this event occurs again.

SCOM alert—Enabled

6230: The job "%1" (ID %2) has started reversing the direction of the protection of %3.

Event log—Application

Source—Double-Take Management Service

Type or Level—Information

User action required—No action required.

SCOM alert—Enabled

6231: The job "%1" (ID %2) has successfully reversed the direction of the protection of %3.

Event log—Application

Source—Double-Take Management Service

Type or Level—Information

User action required—No action required.

SCOM alert—Enabled

6232: The job "%1" (ID %2) has encountered an error reversing the direction of the protection of %3.%n%n

Event log—Application

Source—Double-Take Management Service

Type or Level—Error

User action required—Review the additional error information to identify the problem. Correct the problem and retry the operation. Contact technical support if this event occurs again.

SCOM alert—Enabled

6240: The job "%1" (ID %2) is being deleted.

Event log—Application

Source—Double-Take Management Service

Type or Level—Information

User action required—No action required.

SCOM alert—Enabled

6241: The job "%1" (ID %2) has successfully been deleted.

Event log—Application

Source—Double-Take Management Service

Type or Level—Information

User action required—No action required.

SCOM alert—Enabled

6242: The job "%1" (ID %2) has encountered an error while being deleted.%n%n

Event log—Application

Source—Double-Take Management Service

Type or Level—Error

User action required—Review the additional error information to identify the problem. Correct the problem and retry the operation. Contact technical support if this event occurs again.

SCOM alert—Enabled

6250: The job "%1" (ID %2) protecting %3 has completed its mirror.

Event log—Application

Source—Double-Take Management Service

Type or Level—Information

User action required—No action required.

SCOM alert—Enabled

6300: A failover condition has been met for the full server job "%1" (ID %2).

Event log—Application

Source—Double-Take Management Service

Type or Level—Warning

User action required—Check your source machine and initiate failover, if user intervention for failover is configured. If you bring your source machine back online without initiating failover, the failover condition met state will be canceled.

SCOM alert—Enabled

6500: A cutover condition has been met for the full server migration job "%1" (ID %2).

Event log—Application

Source—Double-Take Management Service

Type or Level—Information

User action required—Initiate cutover.

SCOM alert—Enabled

6700: A cutover condition has been met for the data migration job "%1" (ID %2).

Event log—Application

Source—Double-Take Management Service

Type or Level—Information

User action required—Initiate cutover.

SCOM alert—Enabled

7000: Double-Take Metered Usage is not enabled on server %1.

Event log—Application

Source—Double-Take Management Service

Type or Level—Information

User action required—No action required.

SCOM alert—Enabled

7001: Double-Take Metered Usage could not be enabled on server %1.

Event log—Application

Source—Double-Take Management Service

Type or Level—Error

User action required—Contact your Double-Take service provider.

SCOM alert—Enabled

7002: Double-Take Metered Usage is enabled on server %1. The configured service provider is %2. The configured user name is %3. The configured Metered Usage service address is %4.

Event log—Application

Source—Double-Take Management Service

Type or Level—Information

User action required—No action required.

SCOM alert—Enabled

7003: Double-Take successfully updated the metered license on server %1.

Event log—Application

Source—Double-Take Management Service

Type or Level—Information

User action required—No action required.

SCOM alert—Enabled

7004: Double-Take failed to update the metered license on server %1.

Event log—Application

Source—Double-Take Management Service

Type or Level—Warning

User action required—Confirm the server has Internet access. If you have Internet access and continue to receive this message, contact your Double-Take service provider.

SCOM alert—Enabled

7106: The driver was unable to get valid name information from the Filter Manager for the file %2. (Filename may be truncated.) It cannot be replicated. Please contact technical support.

Event log—System

Source—RepDrv

Type or Level—Error

User action required—Contact technical support.

SCOM alert—Enabled

7107: The driver was unable to get valid name information from the Filter Manager for a file. It cannot be replicated. Please contact technical support.

Event log—System

Source—RepDrv

Type or Level—Error

User action required—Contact technical support.

SCOM alert—Enabled

8100: The driver encountered an unrecoverable internal error. Contact technical support. The last Word in the Data window is the internal error code.

Event log—System

Source—RepDac

Type or Level—Error

User action required—Contact technical support.

SCOM alert—Enabled

8192: Driver failed to allocate Kernel memory. Replication is stopped and server must be rebooted for replication to continue. The last word in the data window is the tag of the allocation that failed.

Event log—System

Source—RepDrv, RepKap, or RepHsm

Type or Level—Error

User action required—Reboot the server and contact technical support if this event occurs again.

SCOM alert—Enabled

8192: Kernel memory is exhausted. Replication is stopped. This may have been caused by low system resources.

Event log—System

Source—RepDrv or RepHsm

Type or Level—Error

User action required—Reboot the server and contact technical support if this event occurs again.

SCOM alert—Enabled

8193: The driver failed to create a thread required for normal operation. This may have been caused by low system resources. Reboot your server and contact technical support if this error occurs again. The last Word in the Data window is the NT status code.

Event log—System

Source—RepDrv

Type or Level—Error

User action required—Reboot the server and contact technical support if this event occurs again.

SCOM alert—Enabled

8196: The maximum amount of memory for replication queuing has been reached. Replication is stopped and memory is being freed.

Event log—System

Source—RepDrv

Type or Level—Warning

User action required—This error is expected when the amount of replication exceeds what can be queued and transmitted on the source server. You do not have to take any action because Double-Take will automatically disconnect, reconnect and remirror (by default) when memory resources are available. However, you may want to consider changes to the source that will reduce the load on the server. See Knowledge Base Article 32410 on the support site for details on the 8196 event and possible steps you can take on your server to help alleviate this condition.

SCOM alert—Enabled

8198: The driver registry path could not be saved. The default registry path will be used.

Event log—System

Source—RepDrv, RepKap, or RepHsm

Type or Level—Warning

User action required—No action required.

SCOM alert—Enabled

8200: The driver failed to allocate a buffer for a file name longer than 260 characters. The file will be skipped. The last Word in the Data window is the NT status code.

Event log—System

Source—RepDrv

Type or Level—Warning

User action required—Reboot the server and contact technical support if this event occurs again.

SCOM alert—Enabled

9000: The driver has failed to process a rename operation. The driver will resend the rename operation. This message is only a warning. If you receive this message repeatedly, contact technical support. The last Word in the Data window is the NT status code.

Event log—System

Source—RepKap

Type or Level—Warning

User action required—Contact technical support if this event occurs again.

SCOM alert—Enabled

9100: The driver encountered an error opening a file from the service. Check the Event Viewer Application log for additional service information or contact technical support. The last Word in the Data window is the exception code.

Event log—System

Source—RepKap

Type or Level—Error

User action required—Check for related service messages. Contact technical support if this event occurs again.

SCOM alert—Enabled

9101: The driver encountered an error reading from the service input buffer. Check the Event Viewer Application log for additional service information or contact technical support. The last Word in the Data window is the exception code.

Event log—System

Source—RepKap

Type or Level—Error

User action required—Check for related service messages. Contact technical support if this event occurs again.

SCOM alert—Enabled

9102: The driver encountered an error writing to the service output buffer. Check the Event Viewer Application log for additional service information or contact technical support. The last Word in the Data window is the exception code.

Event log—System

Source—RepKap

Type or Level—Error

User action required—Check for related service messages. Contact technical support if this event occurs again.

SCOM alert—Enabled

9103: The driver encountered an error writing to the service input buffer. Check the Event Viewer Application log for additional service information or contact technical support. The last Word in the Data window is the exception code.

Event log—System

Source—RepKap

Type or Level—Error

User action required—Check for related service messages. Contact technical support if this event occurs again.

SCOM alert—Enabled

9104: The driver encountered an error querying for file security from the service input buffer. Check the Event Viewer Application log for additional service information or contact technical support. The last Word in the Data window is the exception code.

Event log—System

Source—RepKap

Type or Level—Error

User action required—Check for related service messages. Contact technical support if this event occurs again.

SCOM alert—Enabled

9105: The driver encountered an error querying for file security from the service input buffer. Check the Event Viewer Application log for additional service information or contact technical support. The last Word in the Data window is the exception code.

Event log—System

Source—RepKap

Type or Level—Error

User action required—Check for related service messages. Contact technical support if this event occurs again.

SCOM alert—Enabled

9106: The driver encountered an error writing file security data to the service input buffer. Check the Event Viewer Application log for additional service information or contact technical support. The last Word in the Data window is the exception code.

Event log—System

Source—RepKap

Type or Level—Error

User action required—Check for related service messages. Contact technical support if this event occurs again.

SCOM alert—Enabled

9107: The driver encountered an error querying for an allocated range from the service input buffer. Check the Event Viewer Application log for additional service information or contact technical support. The last Word in the Data window is the exception code.

Event log—System

Source—RepKap

Type or Level—Error

User action required—Check for related service messages. Contact technical support if this event occurs again.

SCOM alert—Enabled

9108: The driver encountered an error querying for an allocated range from the service output buffer. Check the Event Viewer Application log for additional service information or contact technical support. The last Word in the Data window is the exception code.

Event log—System

Source—RepKap

Type or Level—Error

User action required—Check for related service messages. Contact technical support if this event occurs again.

SCOM alert—Enabled

9109: The driver encountered an error writing an allocated range to the service input buffer. Check the Event Viewer Application log for additional service information or contact technical support. The last Word in the Data window is the exception code.

Event log—System

Source—RepKap

Type or Level—Error

User action required—Check for related service messages. Contact technical support if this event occurs again.

SCOM alert—Enabled

9110: The driver encountered an error querying for a directory from the service input buffer. Check the Event Viewer Application log for additional service information or contact technical support. The last Word in the Data window is the exception code.

Event log—System

Source—RepKap

Type or Level—Error

User action required—Check for related service messages. Contact technical support if this event occurs again.

SCOM alert—Enabled

9111: The driver encountered an error querying for a directory from the service output buffer. Check the Event Viewer Application log for additional service information or contact technical support. The last Word in the Data window is the exception code.

Event log—System

Source—RepKap

Type or Level—Error

User action required—Check for related service messages. Contact technical support if this event occurs again.

SCOM alert—Enabled

9112: The driver encountered an error writing a directory query to the service input buffer. Check the Event Viewer Application log for additional service information or contact technical support. The last Word in the Data window is the exception code.

Event log—System

Source—RepKap

Type or Level—Error

User action required—Check for related service messages. Contact technical support if this event occurs again.

SCOM alert—Enabled

9113: The driver encountered an error querying a stream from the service input buffer. Check the Event Viewer Application log for additional service information or contact technical support. The last Word in the Data window is the exception code.

Event log—System

Source—RepKap

Type or Level—Error

User action required—Check for related service messages. Contact technical support if this event occurs again.

SCOM alert—Enabled

9114: The driver encountered an error writing a stream query to the service output buffer. Check the Event Viewer Application log for additional service information or contact technical support. The last Word in the Data window is the exception code.

Event log—System

Source—RepKap

Type or Level—Error

User action required—Check for related service messages. Contact technical support if this event occurs again.

SCOM alert—Enabled

9115: The driver encountered an error writing a stream query to the service output buffer. Check the Event Viewer Application log for additional service information or contact technical support. The last Word in the Data window is the exception code.

Event log—System

Source—RepKap

Type or Level—Error

User action required—Check for related service messages. Contact technical support if this event occurs again.

SCOM alert—Enabled

9116: The driver has failed to close a file handle. If you receive this message repeatedly, contact technical support. The last Word in the Data window is the NT status code.

Event log—System

Source—RepKap

Type or Level—Error

User action required—Contact technical support.

SCOM alert—Enabled

9117: The driver encountered an error querying for extended attributes from the service input buffer. Check the Event Viewer Application log for additional service information or contact technical support. The last Word in the Data window is the exception code.

Event log—System

Source—RepKap

Type or Level—Error

User action required—Check for related service messages. Contact technical support if this event occurs again.

SCOM alert—Enabled

9118: The driver encountered an error writing extended attributes to the service output buffer. Check the Event Viewer Application log for additional service information or contact technical support. The last Word in the Data window is the exception code.

Event log—System

Source—RepKap

Type or Level—Error

User action required—Check for related service messages. Contact technical support if this event occurs again.

SCOM alert—Enabled

9119: The driver encountered an error writing extended attributes status to the service input buffer. Check the Event Viewer Application log for additional service information or contact technical support. The last Word in the Data window is the exception code.

Event log—System

Source—RepKap

Type or Level—Error

User action required—Check for related service messages. Contact technical support if this event occurs again.

SCOM alert—Enabled

9120: The driver encountered an error querying for file information from the service input buffer. Check the Event Viewer Application log for additional service information or contact technical support. The last Word in the Data window is the exception code.

Event log—System

Source—RepKap

Type or Level—Error

User action required—Check for related service messages. Contact technical support if this event occurs again.

SCOM alert—Enabled

9121: The driver encountered an error writing file information to the service output buffer. Check the Event Viewer Application log for additional service information or contact technical support. The last Word in the Data window is the exception code.

Event log—System

Source—RepKap

Type or Level—Error

User action required—Check for related service messages. Contact technical support if this event occurs again.

SCOM alert—Enabled

9122: The driver encountered an error writing file information status to the service input buffer. Check the Event Viewer Application log for additional service information or contact technical support. The last Word in the Data window is the exception code.

Event log—System

Source—RepKap

Type or Level—Error

User action required—Check for related service messages. Contact technical support if this event occurs again.

SCOM alert—Enabled

9123: The driver encountered an error querying for fsctl information from the service input buffer. Check the Event Viewer Application log for additional service information or contact technical support. The last Word in the Data window is the exception code.

Event log—System

Source—RepKap

Type or Level—Error

User action required—Check for related service messages. Contact technical support if this event occurs again.

SCOM alert—Enabled

9124: The driver encountered an error writing fsctl information to the service output buffer. Check the Event Viewer Application log for additional service information or contact technical support. The last Word in the Data window is the exception code.

Event log—System

Source—RepKap

Type or Level—Error

User action required—Check for related service messages. Contact technical support if this event occurs again.

SCOM alert—Enabled

9125: The driver encountered an error writing fsctl status to the service input buffer. Check the Event Viewer Application log for additional service information or contact technical support. The last Word in the Data window is the exception code.

Event log—System

Source—RepKap

Type or Level—Error

User action required—Check for related service messages. Contact technical support if this event occurs again.

SCOM alert—Enabled

9126: The driver encountered an error reading from the service input buffer, KFAI_OPEN_BY_FILE_ID. Check the Event Viewer Application log for additional service information or contact technical support. The last Word in the Data window is the exception code

Event log—System

Source—RepKap

Type or Level—Error

User action required—Check for related service messages. Contact technical support if this event occurs again.

SCOM alert—Enabled

9127: The driver encountered an error writing to the service output buffer, KFAI_OPEN_BY_FILE_ID. Check the Event Viewer Application log for additional service information or contact technical support. The last Word in the Data window is the exception code.

Event log—System

Source—RepKap

Type or Level—Error

User action required—Check for related service messages. Contact technical support if this event occurs again.

SCOM alert—Enabled

9128: The driver encountered an error reading from the service input buffer, KFAI_QUERY_INFO. Check the Event Viewer Application log for additional service information or contact technical support. The last Word in the Data window is the exception code.

Event log—System

Source—RepKap

Type or Level—Error

User action required—Check for related service messages. Contact technical support if this event occurs again.

SCOM alert—Enabled

9129: The driver encountered an error writing to the service output buffer, KFAI_QUERY_INFO. Check the Event Viewer Application log for additional service information or contact technical support. The last Word in the Data window is the exception code.

Event log—System

Source—RepKap

Type or Level—Error

User action required—Check for related service messages. Contact technical support if this event occurs again.

SCOM alert—Enabled

10000: This message is only a placeholder warning. The last Word in the Data window is the NT status code.

Event log—System

Source—Double-Take

Type or Level—Warning

User action required—No action required.

SCOM alert—Enabled

10000: Connect failed to node %1 for resource %2. Adding node to reconnect list.

Event log—Application

Source—Double-Take

Type or Level—Error

User action required—Ensure that GeoCluster is running on all possible owners and that it can communicate on the network selected for mirroring and replication traffic. GeoCluster will try to reestablish a connection using the check unresponsive node interval specified for the resource.

SCOM alert—Enabled

10001: Reconnect succeeded to node %1 for resource %2. Will be added as a possible owner when mirror is complete.

Event log—Application

Source—Double-Take

Type or Level—Information

User action required—No action required.

SCOM alert—Disabled

10002: Disk check failed on node %1 for resource %2. Removing as a possible owner.

Event log—Application

Source—Double-Take

Type or Level—Error

User action required—Ensure that GeoCluster is running on all possible owners and that it can communicate on the public network. Also ensure that the disk specified for the resource is functioning correctly on all possible owners.

SCOM alert—Enabled

10003: Owner %1 of the quorum resource %2 couldn't access the arbitration path %3. Network may be down.

Event log—Application

Source—Double-Take

Type or Level—Error

User action required—Ensure that the network used to access the arbitration path is up and that the server is operational. Also ensure that the arbitration share path does

exist and that the account running the cluster service has write privileges to the share path.

SCOM alert—Enabled

10004: Failover of the group %1 is being delayed. Group will be brought online when the target queue is below the limit or the timeout has expired.

Event log—Application

Source—Double-Take

Type or Level—Warning

User action required—No action required.

SCOM alert—Enabled

10005: Node %1 is taking ownership of the group %2. The group will be brought online on this node.

Event log—Application

Source—Double-Take

Type or Level—Information

User action required—No action required.

SCOM alert—Disabled

10006: The cluster notification thread failed to start on node %1 for resource %2. The resource should be taken offline and brought back online.

Event log—Application

Source—Double-Take

Type or Level—Warning

User action required—Take the resource offline and bring it back online.

SCOM alert—Enabled

10007: The user %1 has reverted a snapshot for the %2 resource on node %3.

Event log—Application

Source—Double-Take

Type or Level—Warning

User action required—No action required. The snapshot you selected will be reverted.

SCOM alert—Enabled

10008: The user %1 has discarded queued data for the %2 resource on node %3.

Event log—Application

Source—Double-Take

Type or Level—Warning

User action required—No action required. The queue you selected will be discarded.

SCOM alert—Enabled

10009: The user %1 is verifying data for the %2 resource on node %3.

Event log—Application

Source—Double-Take

Type or Level—Warning

User action required—A snapshot of the current data has been taken. After you have verified the data, accept or reject the data.

SCOM alert—Enabled

10010: The user %1 has rejected the data for the %2 resource on node %3.

Event log—Application

Source—Double-Take

Type or Level—Warning

User action required—No action required. Since the data was rejected, the data has been reverted to the snapshot taken when the data was selected for verification.

SCOM alert—Enabled

10011: The user %1 has accepted the data for the %2 resource on node %3.

Event log—Application

Source—Double-Take

Type or Level—Warning

User action required—No action required. The current data will be used.

SCOM alert—Enabled

10012: The GeoCluster Replicated Disk resource %1 has been set to validate its data. No data replication is occurring to the remaining nodes in the cluster. Please Accept or Reject the data by right-clicking on the resource and selecting the appropriate option.

Event log—Application

Source—Double-Take

Type or Level—Warning

User action required—Replication has been stopped because of the validation request. Accept or reject the data on the node by right-clicking on the resource and selecting the appropriate option.

SCOM alert—Enabled

10100: The driver could not recall a file because it did not have a token for impersonation. The security provider service should set this token. The last Word in the Data window is the exception code.

Event log—System

Source—RepKap

Type or Level—Error

User action required—Contact technical support if this event occurs again.

SCOM alert—Enabled

10101: The driver could not access the file in the archive bin, due to a failed impersonation attempt. The last Word in the Data window is the exception code.

Event log—System

Source—RepKap

Type or Level—Error

User action required—Contact technical support if this event occurs again.

SCOM alert—Enabled

10102: The driver could not recall the file. The last Word in the Data window is the exception code.

Event log—System

Source—RepKap

Type or Level—Error

User action required—Contact technical support if this event occurs again.

SCOM alert—Enabled

11000: Service has started an archive to %1 (%2) for Replication Set %3, ID: %4

Event log—Application

Source—Double-Take

Type or Level—Information

User action required—No action required.

SCOM alert—Enabled

11001: Service has completed an archive to %1 (%2) for Replication Set %3, ID: %4, %5

Event log—Application

Source—Double-Take

Type or Level—Information

User action required—No action required.

SCOM alert—Enabled

11002: Service has started a recall from %1 (%2) for Replication Set %3, ID: %4

Event log—Application

Source—Double-Take

Type or Level—Information

User action required—No action required.

SCOM alert—Enabled

11003: Service has completed a recall from %1 (%2) for Replication Set %3, ID: %4, %5

Event log—Application

Source—Double-Take

Type or Level—Information

User action required—No action required.

SCOM alert—Enabled

11004: Service has failed connection to the RepHSM driver. %1

Event log—Application

Source—Double-Take

Type or Level—Warning

User action required—Reboot the server or manually restart the RepHSM.sys driver.

SCOM alert—Enabled

11005: Service has aborted the archive operation.

Event log—Application

Source—Double-Take

Type or Level—Warning

User action required—Verify the activation code on the source and target is valid for archiving. Reboot an unlicensed server.

SCOM alert—Enabled

11006: Service has aborted the archive recall operation.

Event log—Application

Source—Double-Take

Type or Level—Warning

User action required—Verify the activation code on the source and target is valid for archiving. Reboot an unlicensed server.

SCOM alert—Enabled

11007: Verification has finished with errors. %1

Event log—Application

Source—Double-Take

Type or Level—Warning

User action required—Review the verification log to correct or accept the errors.

SCOM alert—Enabled

11008: Archive feature is not supported on volume %1

Event log—Application

Source—Double-Take

Type or Level—Warning

User action required—The source and target must be NTFS for archiving functionality.

SCOM alert—Enabled

11009: Service has started an archive preview to %1 (%2) for Replication Set %3, ID: %4

Event log—Application

Source—Double-Take

Type or Level—Information

User action required—No action required.

SCOM alert—Enabled

11010: Service has completed an archive preview to %1 (%2) for Replication Set %3, ID: %4

Event log—Application

Source—Double-Take

Type or Level—Information

User action required—No action required.

SCOM alert—Enabled

11011: Service has aborted the archive preview operation.

Event log—Application

Source—Double-Take

Type or Level—Warning

User action required—Verify the activation code on the source and target is valid for archiving. Reboot an unlicensed server.

SCOM alert—Enabled

12000: The service has started.

Event log—Application

Source—Double-Take

Type or Level—Information

User action required—This message refers to the Double-Take Recall service. No action required.

SCOM alert—Enabled

12001: The service failed to start.

Event log—Application

Source—Double-Take

Type or Level—Error

User action required—Check the user name and password for the Double-Take Recall service to ensure validity. Reinstall the software if this event occurs again.

SCOM alert—Enabled

12002: The service has stopped.

Event log—Application

Source—Double-Take

Type or Level—Information

User action required—This message indicates a system shutdown or the user stopped the Double-Take Recall service. No action is required.

SCOM alert—Enabled

12003: The service failed to create a stop control event. (Error %1)

Event log—Application

Source—Double-Take

Type or Level—Error

User action required—Restart the Double-Take Recall service. Reinstall the software if this event occurs again.

SCOM alert—Enabled

12004: RegisterServiceCtrlHandler failed. (Error %1)

Event log—Application

Source—Double-Take

Type or Level—Error

User action required—Restart the Double-Take Recall service. Reinstall the software if this event occurs again.

SCOM alert—Enabled

12005: Service encountered SetServiceStatus error (Error %1)

Event log—Application

Source—Double-Take

Type or Level—Error

User action required—Restart the Double-Take Recall service. Reinstall the software if this event occurs again.

SCOM alert—Enabled

12006: Service could not get handle to driver for security update. (Error %1)

Event log—Application

Source—Double-Take

Type or Level—Error

User action required—The Double-Take Recall service could not connect to the Double-Take Recall archiving driver. Reboot the server and reinstall the software if this event occurs again.

SCOM alert—Enabled

12007: Service failed a periodic security update. (Error %1)

Event log—Application

Source—Double-Take

Type or Level—Warning

User action required—This message refers to the Double-Take Recall service. The operation will be performed every five minutes. Reinstall the software if this event occurs after five minutes.

SCOM alert—Enabled

12288: The driver encountered an error accessing a buffer from the service. Contact technical support. The last Word in the Data window is the exception code.

Event log—System

Source—RepDrv

Type or Level—Error

User action required—Contact technical support.

SCOM alert—Enabled

16384: The driver encountered an unrecoverable error. Contact technical support.

Event log—System

Source—RepDrv

Type or Level—Error

User action required—Contact technical support.

SCOM alert—Enabled

16385: The driver encountered an unexpected internal result. Contact technical support. The last Word in the Data window is the NT status code.

Event log—System

Source—RepDrv

Type or Level—Error

User action required—Contact technical support.

SCOM alert—Enabled

16393: The driver encountered an internal error. Contact technical support. The last Word in the Data window is the internal error code.

Event log—System

Source—RepDrv

Type or Level—Error

User action required—Contact technical support.

SCOM alert—Enabled

16395: The driver detected a memory error which may have been caused by a bad driver or faulty hardware. Contact technical support. The last Word in the Data window is the internal error code.

Event log—System

Source—RepDrv or RepHsm

Type or Level—Error

User action required—Contact technical support.

SCOM alert—Enabled

16396: The driver failed to create work queues for normal operation. This may have been caused by low system resources. Reboot the server and contact technical support if this error occurs again. The last Word in the Data window is the NT status code.

Event log—System

Source—RepDrv

Type or Level—Error

User action required—Reboot the server and contact technical support if this event occurs again.

SCOM alert—Enabled

16400: RepDrv has encountered an unexpected condition, usually caused by low kernel memory. Unless otherwise mentioned, this event has already been handled and your data remains protected. If you continue to receive these events or have further questions please contact tech support.

Event log—System

Source—RepDrv

Type or Level—Information

User action required—No action required.

SCOM alert—Enabled

Performance Monitor

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

Monitoring Performance Monitor statistics

1. From the Performance Monitor, specify the data to monitor by right-clicking and selecting **Add** or using the **Add** button on the toolbar.
2. Choose one of the following Double-Take Performance Objects.
    - Double-Take Connection
    - Double-Take Kernel
    - Double-Take Security
    - Double-Take Source
    - Double-Take Target
3. Select the statistics you want to monitor, and click **Add**.

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

Performance Monitor should not be used remotely on systems running different operating systems (Windows 2003 to Windows 2008 or vice versa). Performance Monitor can be used remotely when using like systems (Windows 2003 to Windows 2003 or Windows 2008 to Windows 2008).

Performance Monitor statistics

The following tables identify the Double-Take Performance Monitor statistics for each Double-Take counter. For each statistic, you will find the following information.

- **Description**—This description identifies what the statistic is measuring.
- **SCOM alert**—This identifies if a SCOM alert rule for the message is enabled, by default, in the Double-Take Management Pack. If there is no pre-defined rule for this message, that will be indicated. See *Microsoft Systems Center Operations Manager 2007* on page 799 for details on the Management Pack.

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

Double-Take Connection

Bandwidth Limit

Description—The amount of bandwidth that may be used to transfer data

SCOM alert—Disabled

Bytes in disk queue

Description—The number of bytes in the source disk queue

SCOM alert—Disabled

Bytes in replication queue

Description—The number of replication bytes in the source queue

SCOM alert—Disabled

Bytes in the mirror queue

Description—The number of mirror bytes in the source queue

SCOM alert—Disabled

Bytes received

Description—The number of bytes received by the target since the last Performance Monitor refresh

SCOM alert—Disabled

Bytes transferred

Description—The number of bytes transmitted from the source

SCOM alert—Disabled

Compressed bytes transferred

Description—The number of compressed bytes transmitted from the source

SCOM alert—Disabled

Operations in acknowledgement queue

Description—The number of operations waiting in the source acknowledgement queue

SCOM alert—Disabled

Operations in command queue

Description—The number of operations waiting in the source command queue

SCOM alert—Disabled

Operations in mirror queue

Description—The number of mirror operations in the source queue

SCOM alert—Disabled

Operations in replication queue

Description—The number of replication operations in the source queue

SCOM alert—Disabled

Operations received

Description—The number of operations received by the target since the last Performance Monitor refresh

SCOM alert—Disabled

Operations resent

Description—The number of operations re-sent since the last time the Double-Take service was restarted on the source

SCOM alert—Disabled

Operations transmitted

Description—The number of operations transmitted from the source

SCOM alert—Disabled

Task commands queued

Counter—Double-Take Connection

Description—The number of task commands queued on the source

SCOM alert—Disabled

Task commands submitted

Description—The number of task commands submitted on the source

SCOM alert—Disabled

Tasks failed

Description—The number of task commands that have failed to execute on the source

SCOM alert—Disabled

Tasks ignored

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

SCOM alert—Disabled

Tasks succeeded

Description—The number of task commands that have succeeded on the source

SCOM alert—Disabled

Double-Take Kernel

Activation code failures

Description—The number of activation code failures when loading the source or target, since the last time the Double-Take service was restarted on the source

SCOM alert—Disabled

CRC Read Time

Description—The length of time, in microseconds, spent reading CRC (cyclic redundancy check) data on the target. If this value is longer than the standard access time of the target's storage device, it indicates there is possibly an issue reading the data on the target. For example, if the target storage is a SAN, there may be an issue with the way the SAN is configured.

SCOM alert—No rule defined

CRC Thread Count

Description—The number of commands being executed simultaneously on the target. In a properly functioning environment, this number should never be greater than the number of difference mirrors currently being executed on the sources connected to this target. If the value grows larger than the number of currently executing difference mirrors, that indicates there is an error condition.

SCOM alert—No rule defined

Double-Take queue memory usage

Description—The amount of system memory in use by the Double-Take queue

SCOM alert—Disabled

Driver Queue Percent

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

SCOM alert—Disabled

Failed mirror operations

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

SCOM alert—Disabled

Failed replication operations

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

SCOM alert—Disabled

Memory Pool Bytes Available

Description—The amount of memory, in bytes, in the Double-Take memory pool that can be used for Double-Take operations. When Double-Take is at or near idle, the pool bytes available and pool total bytes will at or near equal. If Double-Take is queuing, the pool bytes available will be at or near zero and the pool total bytes will be larger (near 256 MB based on default settings).

SCOM alert—No rule defined

Memory Pool Total Bytes

Description—The amount of memory, in bytes, that Double-Take has allocated for memory pooling. When Double-Take is at or near idle, the pool bytes available and pool total bytes will at or near equal. If Double-Take is queuing, the pool bytes available will be at or near zero and the pool total bytes will be larger (near 256 MB based on default settings).

SCOM alert—No rule defined

Mirror Kbytes generated

Description—The number of mirror kilobytes transmitted to the target. This is the number of bytes generated during mirroring. In other words, this is roughly the amount of traffic being sent across the network that is generated by the mirror. It does not take into account TCP/IP overhead (headers and such), however it does account for attributes and other overhead associated with creating a file. With many small files in a directory, you will see larger statistics than expected because of the file creation overhead. Any subsequent remirror will reset this field to zero and increment from there.

SCOM alert—Disabled

Mirror operations generated

Description—The number of mirror operations transmitted from the source

SCOM alert—Disabled

Open Target Handles

Description—The number of handles currently open on the target.

SCOM alert—Disabled

Replication Kbytes generated

Description—The number of replication kilobytes generated on the source by the file system driver

SCOM alert—Disabled

Replication operations generated

Description—The number of replication operations generated on the source by the file system driver

SCOM alert—Disabled

Double-Take Security

Failed logins

Description—Number of failed login attempts since the last time the Double-Take service was restarted

SCOM alert—Disabled

Successful logins

Description—Number of successful login attempts since the last time the Double-Take service was restarted

SCOM alert—Disabled

Double-Take Source

Auto disconnects

Description—The number of automatic disconnects since the last time the Double-Take service was restarted on the source

SCOM alert—Enabled

Auto reconnects

Description—The number of automatic reconnects since the last time the Double-Take service was restarted on the source

SCOM alert—Disabled

Double-Take Target

Bytes in Disk Queue

Description—The number of bytes in the target disk queue

SCOM alert—Disabled

Bytes in Queue

Description—The number of bytes in the system memory and disk queues

SCOM alert—Disabled

Mirror operations received

Description—The number of mirror operations received on the target

SCOM alert—Disabled

Operations received

Description—The number of operations received on the target

SCOM alert—Disabled

Ops Dropped

Description—The number of operations dropped on the target since the last time the Double-Take service was restarted on the target

SCOM alert—Disabled

Ops Remaining

Description—The number of operations on the target remaining to be applied

SCOM alert—Disabled

Orphan Bytes

Description—The number of orphan bytes removed from the target

SCOM alert—Disabled

Orphan Directories

Description—The number of orphan directories removed from the target

SCOM alert—Disabled

Orphan Files

Description—The number of orphan files removed from the target

SCOM alert—Disabled

Retries

Description—The number of retries performed on the target since the last time the Double-Take service was restarted on the target

SCOM alert—Disabled

Tasks failed

Description—The number of task commands that have failed on the target.

SCOM alert—Disabled

Tasks ignored

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

SCOM alert—Disabled

Tasks succeeded

Description—The number of task commands that have succeeded on the target

SCOM alert—Disabled

Microsoft Systems Center Operations Manager 2007

Microsoft Systems Center Operations Manager 2007 (SCOM) is an enterprise class operations management system that provides event management, proactive monitoring and alerting, reporting and trend analysis, system and application specific knowledge, and configurable task responses to proactively respond to negative trends and alerts. Management Packs are pre-configured collections of these capabilities focused on managing a specific application or hardware type, which can be easily exported and imported into other SCOM environments.

The Double-Take Management Pack was created to help you monitor Double-Take operations and provides the following features.

- **Event rules to monitor all Double-Take generated events**—All Double-Take generated events that appear in the Event Viewer can trigger a SCOM alert. By default, only a subset of events are pre-selected to generate alerts, but additional alerts can easily be generated by enabling additional event rules. See *Event messages* on page 718 for a list of the events.
- **Performance rules for threshold violations**—All Double-Take performance counters that appear in the Performance Monitor can generate a SCOM alert when the configured threshold is violated. Because every environment is unique, only a few performance thresholds are enabled by default. See *Customizing the Management Pack* on page 800 for instructions on configuring the performance thresholds, and see *Performance Monitor statistics* on page 791 for a list of the statistics.
- **Performance rules for performance monitoring**—A subset of Double-Take performance counters can be graphically monitored in the Double-Take performance views, accessible from the SCOM console. These statistics illustrate key metrics, such as how much memory or disk space Double-Take is consuming and how much data is being transmitted.
- **Vendor produced knowledge for alerts**—Understanding why problems exist and how to fix them is an important part of operations management. The Double-Take Management Pack contains product knowledge for each alert, gathered from Vision Solutions technical support. Each alert will also provide information and links to external support.
- **Double-Take specific views**—Various Double-Take specific views are provided in the SCOM console.
    - **Alerts View**—View only Double-Take alerts for computers with Double-Take installed.
    - **State View**—View the server state for all Double-Take servers. You can also view the various properties of all Double-Take servers, including the overall server state.
    - **Events View**—View Double-Take events for Double-Take servers.
    - **Performance Data View**—View graphs of various performance counters for one or multiple computers, as defined in the performance rules.

Installing the Double-Take Management Pack

Microsoft Management Pack for Microsoft Systems Center Operations Manager 2007 R2 is required for the Double-Take Management Pack. To improve the operation of the Double-Take Management Pack, you should have the OpsMgr 2007 MOM 2005 Backward Compatibility MP Update installed. This Management Pack can be found on the Microsoft download site at http://www.microsoft.com/downloads/details.aspx?FamilyID=655cdd06-861e-4342-99b2-8a81e09f6546&DisplayLang=en.

The Management Pack is distributed as an .xml file that is imported via the SCOM console.

1. Locate the Double-Take Management Pack using one of the following methods.
   - Start the Double-Take installation, and when the Autorun appears, select the **Get the SCOM Management Pack** link. Copy the .xml file to your SCOM machine.
   - On the Double-Take installation DVD, browse to the SCMgmt\SCOM directory and copy the .xml file to your SCOM machine.
   - Download the Management Pack from the Vision Solutions support web site to your SCOM machine.
2. From the SCOM console, click **Administration** or select **Go**, **Administration**.
3. Right-click on the **Management Packs** line item in the left pane and select **Import Management Pack**.
4. Navigate to the location of the Doubletake.xml file, and follow the steps in the Management Pack Import Wizard.

See the Microsoft SCOM documentation for complete installation details.

Customizing the Management Pack

After the installation, you will have the following Double-Take Management Pack assets.

- **Double-Take Product Version Attribute**—This attribute checks the Double-Take product version and is used by the Double-Take Servers computer group.
- **Double-Take Servers Group and Installation Type**—This asset uses a formula and regular expression match on the Double-Take Product Version attribute to determine which servers have Double-Take installed.
- **Double-Take Rules**—Rules can be found under the Management Pack Objects section, and contains all of the rules that comprise the Management Pack. Performance Rules, Alert Rules and Event Rules are grouped together under the Rules category. (To see only Double-Take rules, change the scope of the Rules group and filter by Double-Take Servers Installation.)
- **Double-Take views**—Various views are available as previously described in **Double-TakeSpecific Views**.

Except for vendor produced product knowledge, all aspects of the Management Pack can be modified by the SCOM administrator. Use the following notes as suggestions for customization, and see your SCOM documentation for complete instructions.

- **Threshold configuration**—Because each server environment is different, the thresholds at which alerts should be generated will be different. To enable one of the pre-existing (but disabled) performance rules, right-click a performance rule and select **Properties**. On the **General** tab,

enable **Performance Rule**. If you want to customize the threshold trigger, modify the threshold value on the **Overrides** tab.

- **Multi-tier alerting**—Although only one performance rule is provided for each performance counter, by copying that rule and changing the alert severity and threshold values, multi-tiered alerting is possible. For example, it is possible to generate a warning alert when 1024 MB memory is used for queuing and generate an error when 512 MB is consumed. Threshold values can be modified on the **Overrides** tab. Alert severities can be modified on the **Configuration** tab, by highlighting the **GenerateAlert** entry and editing the XML data behind the rule. For more details, see your SCOM documentation.
- **Enable additional event rules**—By default, only a pre-selected group of events will generate alerts. If additional alerts are desired, enable additional Event Rules and verify the alert severity.
- **Notifications**—For each Double-Take Alert Rule, a notification response is pre-selected and will send the applicable message to the defined Recipients and notification Subscriptions. Notification and Subscription options in SCOM do not allow for pre-defined Double-Take notification groups. Therefore, it is required that the SCOM administrator create custom notification Recipients and Subscriptions. (Notification Groups previously used in MOM 2005 can be re-created in SCOM using the Subscriptions feature. For more information on Recipients and Subscriptions, refer to Operations Manager Help using the keyword Notifications.)

SNMP

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

Double-Take installs two components to work with SNMP.

1. The sub-agent is a program that installs and runs on the same machine as Double-Take and gathers statistics, data, and traps. The sub-agent forwards the information to the SNMP agent, which relays the information to the manager. The Double-Take SNMP sub-agent is included in the Double-Take installation program.
2. A Double-Take MIB file is placed on the administrator's machine so that the Management Console can interpret the data sent from the sub-agent. The Double-Take .mib file is dt.mib and meets SNMP standards.

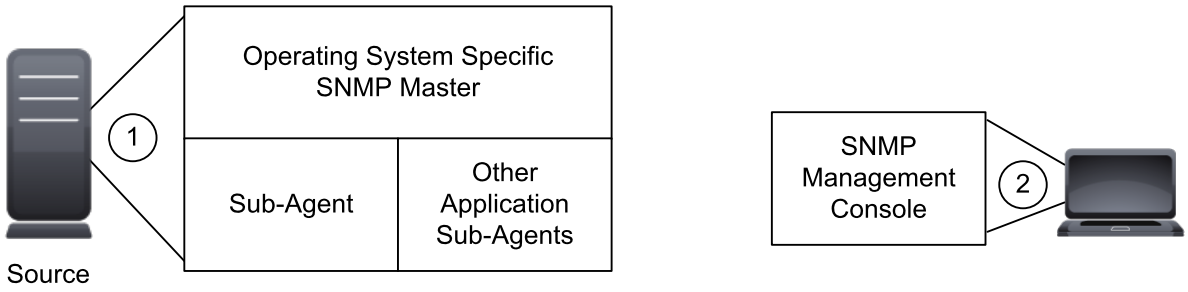

- *Configuring SNMP on your server* on page 802
- *SNMP traps* on page 803
- *SNMP statistics* on page 806

Configuring SNMP on your server

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

Double-Take SNMP will run in any environment, but it only uses 32-bit statistics and traps.

SNMP traps

The following table lists the Double-Take SNMP traps.

---

dttrapAutoDisconnectEndConnection

Auto-disconnect has intentionally dropped the connection

dttrapAutoDisconnectPauseTransmission

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

dttrapAutoDisconnectShutdown

Auto-disconnect forced Double-Take to shut down

dttrapAutoDisconnectWriteQueue

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

dttrapAutoReconnect

Auto-reconnect needs to make a new connection

dttrapConnectionFailed

The source to target connection was not successful

dttrapConnectionLost

The source to target connection has been disconnected

dttrapConnectionPause

The source to target transmission has paused

dttrapConnectionRequested

The source has requested a connection to the target

dttrapConnectionRequestReceived

The target has received a connection request from the source

dttrapConnectionResume

The source to target transmission has resumed

dttrapConnectionSucceeded

The source to target connection has been established

dttrapFailoverConditionMet

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

dttrapFailoverInProgress

Failover or cutover is occurring

dttrapKernelStarted

Double-Take has started

dttrapKernelStopped

Double-Take has stopped

dttrapLicenseViolationOnNetwork

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

dttrapLicenseViolationStartingSource

The source or target cannot be started due to a license violation

dttrapMemoryLimitReached

The Double-Take memory pool limit has been reached

dttrapMemoryLimitRemedied

The memory pool usage is below the maximum limit specified

dttrapMirrorEnd

Mirroring has ended

dttrapMirrorPause

Mirroring has paused

dttrapMirrorResume

Mirroring has resumed

dttrapMirrorStart

Mirroring has started

dttrapMirrorStop

Mirroring has stopped

dttrapReplicationStart

Replication has started

dttrapReplicationStop

Replication has stopped

dttrapRepSetModified

The replication set has been modified

dttrapRestoreComplete

Restoration has ended

dttrapRestoreStarted

Restoration has started

dttrapScheduledConnectEnd

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

dttrapScheduledConnectStart

A scheduled connection has been started

dttrapSourceStarted

The Double-Take source component has started

dttrapSourceStopped

The Double-Take source component has stopped

dttrapTargetFull

The target is full

dttrapTargetStarted

The Double-Take target component has started

dttrapTargetStopped

The Double-Take target component has stopped

dttrapVerificationEnd

Verification has ended

dttrapVerificationFailure

Verification has the source and target are not synchronized

dttrapVerificationStart

Verification has started

SNMP statistics

The following table lists the Double-Take SNMP statistics.

dtActFailCount

The number of activation code errors

dtAutoDisCount

The number of auto-disconnects

dtAutoReCount

The number of auto-reconnects

dtconBytesCompressedTx

The total number of compressed bytes transmitted to the target

dtconBytesInMirQueue

The number of mirror bytes in the queue

dtconBytesInRepQueue

The number of replication bytes in the queue

dtconBytesRx

The total number of bytes received by the target

dtconBytesTx

The total number of bytes transmitted to the target

dtconConnectTime

The length of time, in seconds, that the connection has been active

dtconIpAddress

The IP address of the connected machine. If you are on the source, then this will be the IP address of the target. If you are on the target, then this will be the IP address of the source.

dtConnectionCount

The number of active connections between servers

dtconOpsInAckQueue

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

dtconOpsInCmdQueue

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

dtconOpsInMirQueue

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

dtconOpsInRepQueue

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

dtconOpsRx

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

dtconOpsTx

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

dtconResentOpCount

The number of operations that were resent because of acknowledgement errors

dtconState

The state of the active connection

0—None. This indicates there is no active connection. This may be because the connection has not been established or the underlying connection is unavailable. Statistics are still available for the source and target machines.

1—Active. This indicates that the connection is functioning normally and has no scheduling restrictions imposed on it at this time. (There may be restrictions, but it is currently in a state that allows it to transmit.)

2—Paused. This indicates a connection that has been paused.

4—Scheduled. This indicates a connection that is not currently transmitting due to scheduling restrictions (bandwidth limitations, time frame limitations, and so on).

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

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

dtCurrentMemoryUsage

The amount of memory, in bytes, allocated from the Double-Take memory pool

dtCurrentMemoryUsageMB

The amount of memory, in MB, allocated from the Double-Take memory pool

dtDriverQueuePercent

The percentage of the driver queue that is currently in use. (This is the amount of throttling calculated as a percentage of the stop replicating limit.)

dtFailedLoginCount

The number of unsuccessful logins

dtFailedMirrorCount

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

dtFailedRepCount

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

dtLoginCount

The number of successful logins and logouts

dtMirBytesGenerated

The number of mirror bytes transmitted to the target. This is the number of bytes generated during mirroring. In other words, this is roughly the amount of traffic being sent across the network that is generated by the mirror. It does not take into account TCP/IP overhead (headers and such), however it does account for attributes and other overhead associated with creating a file. With many small files in a directory, you will see larger statistics than expected because of the file creation overhead. Any subsequent remirror will reset this field to zero and increment from there.

dtMirOpsGenerated

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

dtOpsDroppedCount

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

dtRepBytesGenerated

The number of bytes generated by the replication driver

dtRepOpsGenerated

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

dtRetryCount

The number of file operations that have been retried

dtSourceState

0—Source is not running

1—Source is running without the replication driver

2—Source is running with the replication driver

dtTargetState

0—Target is not running

1—Target is running

dtUpTime

The time in seconds since Double-Take was last started

Chapter 19 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

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. Also, Double-Take uses ICMP pings, by default, to monitor the source for failover. You should configure your hardware to allow ICMP pings between the source and target. If you cannot, you will have to configure Double-Take to monitor for a failure using the Double-Take service. See the failover instructions for your specific job type.
- **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.
- **ESX ports**—If you are using VirtualCenter or an ESX host, port 443 is also required and must be opened.

You need to configure your hardware so that the Double-Take ports, Microsoft Windows ports, and ESX 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.

Domain controllers

Failover of domain controllers is dependent on the Double-Take functionality you are using.

- **Domain controller role**—If you want to failover a domain controller, including the roles of the domain controller, you should create full server or virtual protection.
- **Non-domain controller role**—If you are only protecting data, you can failover a domain controller, but the roles of the domain controller are not included. The server will be a member server after failover. In this case, you need to keep in mind that the unavailability of some of the FSMO (Flexible Single Master Operation) roles can cause immediate issues such as the inability to extend the Active Directory schema or to add a domain to a forest.
- **Global catalog server**—If your source is a global catalog server, you should have other global catalog servers throughout the network to ensure that the failure of the source will not impact users.

NetBIOS

Because NetBIOS is not available on Windows 2008 or 2012, if you are using a workgroup environment and do not have DNS host records, there may be a delay of up to five minutes before the failed over source server name is available for client access. See *About the NetBIOS Interface* in the MSDN library.

WINS

When Double-Take failover occurs, Windows initiates WINS registration with the target's primary WINS server to associate the source server's name with the target's primary IP address. In an environment with just one WINS server, no additional processing is required. In an environment with more than one WINS server, WINS replication will distribute the updated WINS registration to other WINS servers on the network. The length of time required for all WINS servers to obtain the new registration depends on the number of WINS servers, the WINS replication architecture, and the WINS replication interval. Clients will be unable to access the target until their WINS server has received the updated WINS information. You can reduce the time required for the WINS updates, thereby decreasing the wait time for the end users, by scripting the WINS updates in the Double-Takefailover scripts. You have two options for scripting the WINS updates.

- *WINS registration* on page 814—This option registers a user-specified server with WINS. It requires less network overhead but administrator group membership on all WINS servers.
- *WINS replication* on page 815—This option forces WINS replication. It does not require any special privileges, but requires system and network resources to complete WINS replication. The impact on the network will depend on the size and complexity of the WINS architecture.

WINS registration

WINS registration can be added to your failover and failback scripts by using the Windows NETSH command with the WINS add name context. Add the following command to your failover and failback scripts for each additional WINS server in your environment (excluding the target's primary WINS server).

netsh wins server ***wins_server_IP_address*** add name Name=***source_server_name*** RecType=1 IP={***IP_address***}

Use the following variable substitutions.

- ***wins_server_IP_address***—The IP address of the WINS server
- ***source_server_name***—The name of the source server
- ***IP_address***—The IP address of the target that has taken over for the failed source (for the failover script) or the IP address of the source that is reassuming its original identity (for the failback script)

For example, suppose you had the following environment.

- **Source name and IP address**—Alpha 192.168.1.108
- **Target name and IP address**—Beta 116.123.2.47
- **Target's Primary WINS server**—116.123.2.50
- **First secondary WINS server on the network**—192.168.1.110
- **Second secondary WINS server on the network**—150.172.114.74

You would add the following to your failover script to register the source's name with the target's IP address on the two secondary WINS servers.

```
netsh wins server 192.168.1.110 add name Name=Alpha RecType=1 IP={116.123.2.47}
netsh wins server 150.172.114.74 add name Name=Alpha RecType=1 IP={116.123.2.47}
```

You would add the following to your failback script to register the source's name back with the source's original IP address on the two secondary WINS servers.

```
netsh wins server 192.168.1.110 add name Name=Alpha RecType=1 IP={192.168.1.108}
netsh wins server 150.172.114.74 add name Name=Alpha RecType=1 IP={192.168.1.108}
```

See your Windows documentation or the Microsoft web site for more details on the NETSH command.

WINS replication

WINS replication can be added to your failover and failback scripts by using the Windows NETSH command with the WINS set replicate context. Add the following command to your failover and failback scripts.

netsh wins server ***target's_primary_wins_server_IP_address*** set replicateflag 1

Use the following variable substitution.

- ***target's_primary_wins_server_IP_address***—The IP address of the target's primary WINS server

For example, suppose you had the following environment.

- **Source name and IP address**—Alpha 192.168.1.108
- **Target name and IP address**—Beta 116.123.2.47
- **Target's Primary WINS server**—116.123.2.50
- **First secondary WINS server on the network**—192.168.1.110
- **Second secondary WINS server on the network**—150.172.114.74

You would add the following to your failover script to force the target's primary WINS server to replicate its updated information to the other secondary WINS servers on the network.

```
netsh wins server 116.123.2.50 set replicateflag 1
```

You would add the same line to your failback script to force the target's primary WINS server to replicate its updated information again. This would replicate information for the source's name and the source's original IP address to the other secondary WINS servers on the network.

```
netsh wins server 116.123.2.50 set replicateflag 1
```

See your Windows documentation or the Microsoft web site for more details on the NETSH command.

DNS

If you are using a Microsoft DNS server, when Double-Take failover occurs, DNS may or may not be automatically updated depending on your job type and job options. If the end-users use DNS to resolve server names and the source IP address was not failed over to the target, additional DNS updates will be required because the host records for the source will remain intact after failover. You can automate this process by scripting the DNS updates in the failover and failback scripts. You have two options for scripting the DNS updates.

- *Windows DNSCMD command* on page 816—The Windows Support Tools contain a DNS Server Troubleshooting Tool utility. This utility includes the DNSCMD command which can be scripted to delete and add host and reverse lookup entries in DNS.
- *Double-Take DFO utility* on page 818—Double-Take also has a utility, called DFO (DNS Failover). The DFO utility can be used to script the deletion and addition of the host and reverse lookup entries in DNS. This utility is installed with Double-Take.

Windows DNSCMD command

DNS updates can be added to your failover and failback scripts by using the Windows DNSCMD command as long as dynamic updates are enabled on the DNS zone and the account running the Double-Take service is a member of the DNSAdmins security group. (See your Microsoft documentation to verify if dynamic updates are enabled.) You may want to disable the DNS registration feature of each IP address that is being changed in DNS to prevent the source from changing the record back when it comes online after a failover.

Add the following commands to your failover and failback scripts to delete the host and reverse lookup entries and add new entries associating the source to the target.

- dnscmd DNS_server's_FQDN /RecordDelete DNS_zone source_server_name A source_server_IP_address /f
- dnscmd DNS_server's_FQDN /RecordDelete www.xxx.in-addr.arpa zzz.yyy PTR source_server's_FQDN /f
- dnscmd DNS_server's_FQDN /RecordAdd DNS_zone source_server_name A target_server_IP_address
- dnscmd DNS_server's_FQDN /RecordAdd aaa.bbb.in-addr.arpa ddd.ccc PTR source_server's_FQDN

Use the following variable substitutions.

- DNS_server's_FQDN—The fully qualified domain name of the DNS server
- DNS_zone—The name of the DNS zone
- source_server_name—The name of the source server
- source_server_IP_address—The IP address on the source
- www.xxx—The first two octets of the source's IP address. For example, if the source's IP address is 192.168.1.108, this variable would be 192.168.
- zzz.yyy—The last two octets, in reverse order, of the source's IP address. For example, if the source's IP address is 192.168.1.108, this variable would be 108.1.
- source_server's_FQDN—The fully qualified domain name of the source server

- target_server_IP_address—The IP address on the target
- aaa.bbb—The first two octets of the target's IP address. For example, if the target's IP address is 116.123.2.47, this variable would be 116.123.
- ddd.ccc—The last two octets, in reverse order, of the target's IP address. For example, if the target's IP address is 116.123.2.47, this variable would be 47.2.

For example, suppose you had the following environment.

- Full qualified domain name of the source—Alpha.domain.com
- Source IP address—192.168.1.108
- Fully qualified domain name of the target—Beta.domain.com
- Target IP address—116.123.2.47
- Fully qualified domain name of the DNS server—DNSServer.domain.com
- DNS zone—domain.com

You would add the following to your failover script to delete the host and reverse lookup entries and add new entries associating the source to the target.

```
dnscmd DNSServer.domain.com /RecordDelete domain.com alpha A 192.168.1.108 /f
dnscmd DNSServer.domain.com /RecordDelete 192.168.in-addr.arpa 108.1 PTR alpha.domain.com /f
dnscmd DNSServer.domain.com /RecordAdd domain.com alpha A 116.123.2.47
dnscmd DNSServer.domain.com /RecordAdd 116.123.in-addr.arpa 47.2 PTR alpha.domain.com
```

You would add the following to your failback script to delete the host and reverse lookup entries and add new entries associating the source with its original identity.

```
dnscmd DNSServer.domain.com /RecordDelete domain.com alpha A 116.123.2.47 /f
dnscmd DNSServer.domain.com /RecordDelete 116.123.in-addr.arpa 47.2 PTR alpha.domain.com /f
dnscmd DNSServer.domain.com /RecordAdd domain.com alpha A 192.168.1.108
dnscmd DNSServer.domain.com /RecordAdd 192.168.in-addr.arpa 108.1 PTR alpha.domain.com
```

See your Windows documentation or the Microsoft web site for more details on the DNSCMD command.

Double-Take DFO utility

DNS updates can be added to your failover and failback scripts by using the Double-Take DFO utility as long as the utility has been registered and the proper privileges are configured.

How the DFO utility works

The DFO utility performs DNS resource record modifications by connecting to the DNS namespace (root\microsoftdns) on the DNS server using WMI. The WMI connection can be made using passed credentials or impersonation if the account running the DFO utility has permissions to perform all DNS-related activities. Passed credentials can be encrypted using Microsoft's CAPICOM dynamic link library with DFO specifying the triple DES encryption algorithm with the maximum key length available (168). By providing reliable encryption, the DFO utility allows you to avoid storing secure passwords in script files.

If the source experiences a failure or an extended outage, clients will need to be redirected automatically to the target server. In these cases, the DFO utility can help make the network redirection portion of failover transparent to end users.

The DFO utility is able to modify five DNS resource record types: A, AAAA, CNAME, MX, and PTR. Here is how it works for the host record or A type.

1. The DFO utility builds and executes a focused WMI query to retrieve a collection of matching source DNS resource records from the DNS server. For example: SELECT * FROM MicrosoftDNS_AType WHERE IPAddress="192.168.1.108"
2. The DFO utility iterates through the returned collection and modifies any matching resource records.
    a. The DFO utility spawns an instance of the WMI DNS resource record object to call the modify method.
    b. The DFO utility sets the parameters such that the target IP address is the input parameter.
    c. The DFO utility executes the modify method on the WMI DNS record object.
3. The DFO utility locks the DNS resource record in Active Directory so that the source computer account is unable to modify the record outside of the DFO utility. This is done to prevent modification of the resource record to point back to the source machine until the failback process is initiated by the user through the DFO utility.
    a. The DFO utility denies permission to modify the Active Directory object representing the DNS entry.
    b. The DFO utility gets the DNS resource record object in Active Directory.
    c. The DFO utility reads the security descriptor and gets the DACL.
    d. The DFO utility adds the ACE "ACCESS_DENIED" type for the passed-in trustee name (for example, source computer account, cluster administrator account, and so on) to deny access to the "Write All Properties" permission.
4. The DFO utility logs the results of the actions performed/attempted.

Other record types require different queries and input parameters. Additionally, CNAME, MX, and PTR record types do not execute the Active Directory object locking routines that A and AAAA type records require for failover.

- **AAAA type**—Except for the query difference, this record type is identical to the A type record.

  SELECT * FROM MicrosoftDNS_AAAAType WHERE IPAddress="21DA:D3:0:2F3B:2AA:FF:FE28:9C5A"

- **CNAME type**—This type does not have Active Directory object locking to prevent updates during failover.

  SELECT * from MicrosoftDNS_CNAMEType WHERE PrimaryName="sql1.doubletake.com"

- **MX type**—This type does not have Active Directory object locking to prevent updates during failover.

  SELECT * from MicrosoftDNS_MXType WHERE MailExchange="mail1.doubletake.com"

- **PTR type**—Instead of modifying the source record, the PTR type deletes the source PTR record and create a new PTR record by using previous source PTR text record information, substituting the target FQDN for the source FQDN, and calling the CreateInstanceFromPropertyData() method on the DNS server. This type does not have Active Directory object locking to prevent updates during failover.

  SELECT * from MicrosoftDNS_PTRType WHERE PTRDomainName="sql1.doubletake.com"

During failback, the same mechanisms that were used during failover are used, except that the original source-related records are modified to point to the original source. (During failover, the source records were modified to point to the target IP address or name, depending on the record type.) Also, during failover the A and AAAA type DNS resource records are modified in DNS and then locked in Active Directory; during failback, those record types are unlocked in Active Directory and then modified in DNS.

Using the DFO utility

1. From a command prompt, change to the Double-Take program files directory and register the DFO utility by entering the command regsvr32 capicom.dll
2. Create a user account that has full control on the WMI DNS namespace on the source's primary DNS server.
   a. If you are using Windows 2003 SP1, use the following steps.
      1. From a command prompt, enter the command mmc.
      2. After the Microsoft Management Console starts, select **File**, **Add/Remove Snap-in**.
      3. Click **Add**, select **WMI Control**, click **Add** again, confirm the local computer is selected, and then click **Finish**.
      4. Close the snap-in dialog box and then click **OK** to return to the console.
   b. If you are using Windows 2003 SP2 or later, from a command prompt, enter the command wmimgmt.msc.
   c. Right-click **WMI Control** and select **Properties**.
   d. On the **Security** tab, expand the tree under **Root**.
   e. Select **MicrosoftDNS** and click **Security**.
   f. Click **Add** and identify the user account that you want the DFO utility to use.
   g. If you are using Windows 2003 SP1, grant the user account permissions for Execute Methods, Full Write, Partial Write, Provider Write, Enable Account, Remote Enable, and Read Security.

h. If you are using Windows 2003 SP2 or later, use the following steps.
    1. Grant the user account permissions for Execute Methods, Enable Account, Remote Enable, and Read Security.
    2. Click **Advanced** and in the **Permissions** list, select the user account and click **Edit**. Select **This namespace and subnamespaces**.
i. Click **OK** to close all open dialog boxes and then close the console.
j. Restart the Windows Management Instrumentation service for the changes to take effect. .

3. If you are using Windows 2003 SP2 or later, complete the following additional steps.
    a. From a command prompt, enter the command dcomcnfg.
    b. Expand **Component Services**, expand **Computers**, then right-click **My Computer** and select **Properties**.
    c. On the **COM Security** tab, under **Access Permissions**, click **Edit Limits**.
    d. Click **Add**, identify the user account, and click **OK**.
    e. In the **Permissions for User** list, allow permissions for Local Access and Remote Access and click **OK**.
    f. Under **Launch and Activation Permissions**, click **Edit Limits**.
    g. Click **Add**, identify the user account, and click **OK**.
    h. In the **Permissions for User** list, allow permissions for Local Launch, Remote Launch, Local Activation, and Remove Activation.
    i. Click **OK**.
    j. Click **OK** again.
    k. Expand **My Computer**, expand **DCOM Config**, then right-click **Windows Management and Instrumentation** and select **Properties**.
    l. On the **Security** tab, under **Access Permissions**, click **Edit**.
    m. Click **Add**, identify the user account, and click **OK**.
    n. In the **Permissions for User** list, allow permissions for Local Access and Remote Access and click **OK**.
    o. Click **OK** to close all open dialog boxes.
    p. Restart the DNS/Domain Controller.
4. Add the same user account that has full control on the WMI DNS namespace to the domain's DnsAdmins group where the source's primary DNS server is located.
    a. Select **Active Directory Users and Computers** from Administrative Tools.
    b. Right-click the **DnsAdmins** group and select **Properties**.
    c. Select the **Members** tab, click **Add**, and identify the user account that you granted full control on the WMI DNS namespace.
    d. Click **OK** to close all open dialog boxes and then close Active Directory Users and Computers.
5. Add a user to the Server Operator group.
    a. Select **Active Directory Users and Computers** from Administrative Tools.
    b. Select **Builtin**, then right-click the **Server Operators** group and select **Properties**.
    c. Select the **Members** tab, click **Add**, and identify the user account that you granted full control on the WMI DNS namespace.

d. Click **OK** to close all open dialog boxes and then close Active Directory Users and Computers.

6. Grant the user full control over the source and target DNS records.
   a. Select **DNS** from Administrative Tools.
   b. Locate both the source and target records in the forward and reverse lookup zones.
   c. For each record, right-click and select **Properties**.
   d. On the **Security** tab, click **Add** and identify the user account that you granted full control on the WMI DNS namespace, and click **OK**.
   e. In the **Permissions for User** list, allow permissions for **Full control** and click **OK**.
   f. Click **OK** to close all open dialog boxes and repeat for each record.
   g. Close DNS Manager.
7. Add the appropriate DFO command to your failover script using the following syntax.

Command

DFO

Description

Used in scripts to failover DNS server name

Syntax

DFO [/DNSSRVNAME <dns_server_name>] [/SRCNAME <source_fqd_name>] [/SRCIP <source_ip>] [/TARIP <target_ip>] [/TARNAME <target_fqd_name>] [/RECORDTYPE <rec_type>] [/USERNAME <user_name>] [/PASSWORD <password>] [/DNSZONE <zone_name>] [/DNSDOMAIN <domain_name>] [/LOGFILE <file_name>] [/FAILBACK [fb_switch]] [/SETPASSWORD <user_name> <password>[machine][file]] [/GETPASSWORD] [/LOCK] [/UNLOCK] [/TRUSTEE [<trustee_name>]] [/VERBOSE] [/FLUSHDNS] [/MACHINE <machine_fqd_name>] [/TTL <seconds>] [/ADDOMAIN <active_directory_domain_name>] [/SOURCEDN <source_domain_name>] [/TEST] [/DEBUG] [/HELP]

Options

- DNSSRVNAME dns_server_name—The name of the source domain/zone's primary DNS server. If not specified, the local machine will be used.
- SRCNAME source_fqd_name—The source machine's fully qualified domain name
- SRCIP source_ip—The source machine's IP address
- TARIP target_ip—The target machine's IP address
- TARNAME target_fqd_name—The target machine's fully qualified domain name (required only for failback)
- RECORDTYPE rec_type—The type of DNS resource records to modify or list. Values record types are ALL, MSEXCHANGE, A, CNAME, MX, PTR, STD, or STANDARD. STD and STANDARD are used to specify a non-Exchange record (minus the MX records). By default, all record types are included.

- USERNAME user_name—The domain name of the user account. If not specified, the account running the utility will be used.
- PASSWORD password—The password associated with the user account
- DNSZONE zone_name—The name of the DNS zone or DNS container, used to refine queries
- DNSDOMAIN domain_name—The name of the DNS domain, used to refine queries
- LOGFILE file_name—The name of the log file
- FAILBACK fb_switch—Denotes a failback procedure, performed after a failed source is recovered or restored (required for failback). By default, the DFO will only failback records in the dfo_failback_config.dat file. The fb_ switch is optional and allows you to enter search criteria to identify the records to change back, even if they are not in the configuration file. The fb_switch is also used if the dfo_failback_config.dat file is missing.
- SETPASSWORD user_name password machine file—Stores user credentials on the specified machine or in the specified file for later use. The file will be encrypted. This option must be run separately from a modify or list activity.
- GETPASSWORD—Retrieves previously stored user credentials. This option can only be used if the credentials were previously stored with the setpassword option.
- LOCK—Allows Active Directory locking for the A record type of the source specified without modifying the record
- UNLOCK—Allows Active Directory unlocking for the A record type of the source specified without modifying the record
- TRUSTEE trustee_name—The domain account for the source machine (domain\machine$). DFO attempts to deny write permissions to the DNS A record on failover for the account identified as the trustee. “Deny write permissions” is then removed from the DNS A record on failback. This keeps the source server from reclaiming its DNS A record if it comes back online prior to failback.
- VERBOSE—Logging and display level set to maximum detail
- FLUSHDNS—Runs the ipconfig /flushdns command to flush the DNS cache.
- MACHINE machine_fqd_name—Specifies the machine where ipconfig /flushdns is run. Use the fully-qualified domain name of the machine.
- TTL seconds—Specifies the number of seconds for the time to live value of all modified records
- ADDOMAIN active_directory_domain_name—The name of the Active Directory domain
- SOURCEDN source_domain_name—The name of the source's domain
- TEST—Runs in test mode so that modifications are not made, only listed
- DEBUG—Forces DFO to write the DNS resource record as-is to the dfolog.log file prior to any DFO modify or list activity.
- HELP—Displays the syntax of the DNS Failover utility

Examples

- dfo /dnssrvname gamma.domain.com /srcname alpha.domain.com /srcip 206.31.4.10 /verbose (Lists all resource records on the specified DNS server that match the source criteria)
- dfo /dnssrvname gamma.domain.com /srcname alpha.domain.com /srcip 206.31.4.10 /tarip 210.11.12.13 /verbose (Modifies all resource records on the specified DNS server that match the source criteria, using the credentials of the account running the utility to connect to the DNS server)
- dfo /dnssrvname gamma.domain.com /srcname alpha.domain.com /srcip 206.31.4.10 /tarip 210.11.12.13 /username domain.com\admin /password /verbose (Modifies all resource records on the specified DNS server that match the source criteria, using the username and password to connect to the DNS server)
- dfo /dnssrvname gamma.domain.com /srcname alpha.domain.com /srcip 210.11.12.13 /tarname beta.domain.com /tarip 206.31.4.10 /failback /verbose (Fails back all resource records on the specified DNS server that were changed on failover)
- dfo /setpassword domain.com\admin password (Stores the user name and password in an encrypted file)
- dfo /dnssrvname gamma.domain.com /srcname alpha.domain.com /srcip 206.31.4.10 /tarip 210.11.12.13 /username domain.com\admin /getpassword /verbose (Modifies all resource records on the specified DNS server that match the source criteria, using the specified username and retrieving the password from the encrypted file)

Notes

All options are marked as optional, enclosed in brackets [ ], however, you will have to supply options to execute DFO functionality. The options to supply will depend on the functionality you are trying to complete. For example, you must supply the username and password to cache credentials, but you do not need those options to query or modify a DNS record.

Non-Microsoft DNS

If you are using a non-Microsoft DNS server (such as Unix) or if you are in a non-domain configuration (such as a workgroup), when Double-Take failover occurs, DNS may or may not be automatically updated depending on your job type and your job options. If the end-users use DNS to resolve server names and the source IP address was not failed over to the target, additional DNS updates will be required because the host records for the source will remain intact after failover. You can automate this process by scripting the DNS updates in the failover and failback scripts.

One option is to use a BIND DNS client for DNS scripting. The following steps provide an example of how you can use a BIND DNS client for DNS failover and failback scripting. You may need to modify this example to fit your environment.

1. Go to www.isc.org and download the appropriate BIND DNS client.
2. Install the BIND client on the target server.
3. Set a PATH statement for the BIND directory to ensure that it runs every time the executable is called.
4. Create a failover script file in the Double-Take directory.
5. Add the following line to the failover script file, substituting your Double-Take directory for install_location.

```
nsupdate.exe "c:\install_location\dnsover.txt"
```

6. Save the failover script file.
7. Create a text file, called dnsover.txt in the Double-Take directory.
8. Add the following lines to the dnsover.txt file, substituting your source name, fully-qualified domain name, target name, and target IP address as appropriate.

```
update delete source_server_name.fully_qualified_domain_name.com A
update add target_server_name.fully_qualified_domain_name.com 86400 A target_server_IP_address
send
```

9. Save the dnsover.txt file.
10. Create a failback script file in the Double-Take directory.
11. Add the following line to the failback script file, substituting your Double-Take directory for install_location.

```
nsupdate.exe "c:\install_location\dnsback.txt"
```

12. Save the failback script file.
13. Create a text file, called dnsback.txt in the Double-Take directory.
14. Add the following lines to the dnsback.txt file, substituting your target name, fully-qualified domain name, source name, and source IP address as appropriate.

```
update delete target_server_name.fully_qualified_domain_name.com A
update add source_server_name.fully_qualified_domain_name.com 86400 A source_server_IP_address
send
```

15. Save the dnsback.txt file.

16. Change the Double-Take service on the target server to a domain account that has rights to modify BIND DNS. Stop and start the service to have it take effect.

Macintosh shares

A share is any volume, drive, or directory resource that is shared across a network. During failover, the target can assume or add any source shares so that they remain accessible to the end users. Automatic share failover only occurs for standard Windows file system shares. Other shares, including Macintosh volumes, must be configured for failover through the failover scripts or created manually on the target.

1. On your target, set the File Server for Macintosh service to manual startup. This allows the post-failover script on the target to control when the service starts on the target.
2. Create each volume on the target machine exactly as it exists on the source. Use the Shared Folder wizard to configure each volume as a Macintosh-accessible volume. Follow these steps to start the wizard.
    a. Open the Control Panel and click **Administrative Tools**.
    b. Select **Configure Your Server**.
    c. In the Configure Your Server window, click the **File Server** link.
    d. Click **Start the Shared Folder wizard** to start the wizard, and then follow the directions provided by the wizard. On the Create Shared Folders screen, you must enable **Apple Macintosh**.

    

    You can automate the creation of the volumes during the failover process by using the macfile volume command in the post-failover batch file. For detailed information on how to use this command, see your Windows reference guide.

3. On the target machine, copy the chngname utility, chngname.exe, from the \tools\Win2K directory of the Double-Take DVD or from the Vision Solutions support web site to the directory where Double-Take is installed.
4. Add the following to your failover script.

```
rem Commands for Macintosh-accessible volume failover
rem The chngname utility (chngname.exe) must be located in the same
rem directory where Double-Take is installed.
rem The following command temporarily changes the name of the server. You
rem will need to replace <drive>:\<directory>\ with the location of
rem your Double-Take chngname utility and replace
rem source_name with the name of the source machine.
<drive>\<directory>\chngname /s source_name
rem The following command starts the File Server for Macintosh service
net start "File server for Macintosh"
rem The following command changes the name of the server back to its
rem original name. You will need to replace <drive>:\<directory>\ with
rem the location of your Double-Take chngname utility.
<drive>\<directory>\chngname /t
```

In the event of a failure, the Macintosh clients must remap the volume in order to access it. From the Macintosh client, use the Chooser to select the volume that needs to be remapped.

NFS Shares

A share is any volume, drive, or directory resource that is shared across a network. During failover, the target can assume or add any source shares so that they remain accessible to the end users. Automatic share failover only occurs for standard Windows file system shares. Other shares, including NFS shares, must be configured for failover through the failover scripts or created manually on the target.

1. On your target, set the NFS service to manual startup. This allows the post-failover script on the target to control when the service starts on the target.
2. Create each shared drive or directory on the target exactly as it exists on the source. Configure each drive or directory as an NFS share by following these steps.
    a. Right-click the drive or directory that you want to share, select **Sharing**, and click the **NFS Sharing** tab on the Program Files Properties dialog box.
    b. Enable **Share this folder**, provide the name of the share, and click **OK**.
3. On the target machine, copy the chngname utility, chngname.exe, from the \tools\Win2K directory of the Double-Take DVD or from the support web site to the directory where Double-Take is installed.
4. Add the following to your failover script.

```
rem Commands for NFS share failover
rem The chngname utility (chngname.exe) must be located in the same
rem directory where Double-Take is installed.
rem The following command temporarily changes the name of the server. You
rem will need to replace <drive>:\<directory>\ with the location of
rem your Double-Take chngname utility and replace
rem source_name with the name of the source machine.
<drive>\<directory>\chngname /s source_name
rem The following command starts the NFS service
net start "Server for NFS"
```

In the event of a failure, the clients must remount the shares in order to access them.

Chapter 20 Recommended optimizations

Double-Take is an exceptionally flexible product that can be used in a wide variety of network configurations. However, this flexibility can make implementing Double-Take effectively difficult. There is often a balance that must be found between various configuration options and their relative benefits.

Through years of testing and implementing in diverse environments, Vision Solutions has compiled the following list of recommended optimizations. Keep in mind, what works for one environment or configuration may not work in another. A best practice in one organization may be ineffective in another. You should work with Vision Solutions technical support or Professional Services when making optimization changes.

Planning

Before you begin your Double-Take installation, you should plan your implementation strategy. Ask yourself the following questions.

- What is the role of each server? Will this server be a source? Will this server be a target?
- Is the source server a Domain Controller? Or does it have another very specific role or configuration? You may want consider protecting the entire server in these cases.
- Is the source running Microsoft Exchange or Microsoft SQL?
- How much data will you be protecting? Can your target handle that amount of data?
- How much bandwidth is available between your source and target? Can your network handle the mirroring and replication traffic between the two servers? If the amount of change is greater than the bandwidth available, you may want to consider getting additional bandwidth or planning for disk queuing.

If there are concerns about resource utilization or how Double-Take replication will impact the environment, you can profile the source server, the network links between the source and target, and the target server before installing Double-Take to ensure that each component has adequate resources to handle the added load of replicating the data. Most environments do not require this type of analysis, but it may be needed if there are applications producing high-volume file writes or limited CPU, memory, disk, or network resources.

The best way to understand the impact of replication in an environment is to set up test equipment that simulates the production environment. However, if the resources to test in this manner are not available, resource utilization can be analyzed using Windows Performance Monitor and a utility to monitor network utilization. Performance data should be logged for a period that encompasses normal usage as well as any maintenance, backup, scheduled jobs, or batch processing that occurs. If utilization of any component is extremely high for a significant period of time, then it may be necessary to modify particular Double-Take options. Keep in mind that some factors that are typically not in a test environment, such as backups and other applications using bandwidth, can affect resource utilization in the production environment.

One method to avoid for planning purposes is estimating the amount of data that will be replicated in a given period using the amount of data backed up in a differential backup. Although this may be valid in some cases, it is usually not a good indicator because it is based on the differences in data at the time of backup. For example, if a 1 MB Microsoft Word document is saved ten times throughout the day, this will result in 10 MB of replication traffic because Word rewrites the entire file each time it is saved. However, this will only result in 1 MB being backed up for a differential backup.

Installation optimizations

Make sure you review all of the *Core Double-Take requirements* on page 23 and any requirements that are specific to your job type. When you perform the installation, you will have several decisions to make.

- **Login**—Always log on to the server with an account that is in the local Administrators group before starting the installation.
- **Components**—Decide what components to install and where to install them. Keep in mind that server components are required for systems that will function as a source or target, and they require an activation code for the service to run. Client components do not require an activation code, but are required to administer Double-Take servers throughout your environment.
- **Activation code**—The activation code that is required for each server is a 24-character, alpha-numeric code which applies the appropriate license to your installation.
- **Queues**—The installation will prompt you to select disk queue settings. Double-Take uses system memory to store data. When the Double-Take system memory limit is reached, Double-Take will queue to disk.
    - If you set the system memory limit 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. In general, the amount of memory Double-Take and other applications on the server are configured to use should be less than the amount of physical memory on the system to prevent low memory conditions.
    - Select your disk queue location for optimal performance. For example, do not put it on the same physical device as the data being replicated. If possible, put it on a dedicated array optimized for writing. If you expect large amounts of disk queuing, you may want to increase the size of the queue files from the default of 5 MB to 50 MB for more efficient queuing. See *QJournalFileSize* on page 136.

    See *Double-Take queue* on page 88 for more details on the disk queue usage.
- **Upgrades**—Keep the following caveats in mind when upgrading.
    - If Double-Take does not function correctly after the upgrade, run the Double-Take Setup, select the **Repair** option, and reboot the server. If Double-Take does not function correctly after the repair, uninstall Double-Take, reboot, and install the new version.
    - If your current Double-Take version is more than two minor versions old, you may want to consider uninstalling the old version and installing the new version instead of upgrading.
    - Always upgrade the target server first when upgrading a source and target configuration.

General optimizations

The following are general optimizations that can be used for any Double-Take job type.

Performance optimizations

- **Initial mirror across slow network**—A large amount of data that is being mirrored across a slow network may take days to complete the initial mirror, depending on the amount of data and the available bandwidth. You may want to consider the following options to reduce the amount of time for the initial mirror.
    - Move the target server to the source's site for the initial mirror. When the mirror is complete, delete the job, move the target server to its permanent location and create a new job using a difference mirror.
    - Archive the data to media that can be transported to the target site and restored to the target server. When the data is restored to the target, create the job using a difference mirror.
    - Create a compressed archive of the source data, copy the archive to the target, decompress the data, and then create the job using a difference mirror.
- **Compression**—Double-Take compression should be used when network bandwidth between the source and target is limited. In some cases, performance may also be improved by enabling compression in high-bandwidth environments. The best level of compression for a given solution will depend on a number of factors, including the type of data being replicated, CPU load, and available bandwidth. Since compression settings can be changed dynamically, the easiest way to find the best level of compression is to enable the mid-level and monitor the results. If data is still being queued on the source , increase the compression level. If CPU load becomes an issue on the server, decrease the compression level.
- **Low bandwidth and queuing**—In low bandwidth environments, you may need to revisit the queuing configuration you established when you were installing Double-Take. See the *Installation optimizations* on page 830 and *Double-Take queue* on page 88 for more details on the disk queue usage.
- **High latency and mirror packet size**—In a high latency environment (greater than 100 ms response times), you may want to consider increasing the size of the packets of mirror data. The default value is 65536 bytes. You may want to double that to 131072 bytes. However, if the average size of the files on the source is smaller than the value you set, changing the value will not help. This option is available through the *Source server properties* on page 92.
- **High latency and MaxChecksumBlocks**—In a high latency environment (greater than 100 ms response times), you may want to consider increasing the number of checksum values retrieved from the target. The default is 32. You may want to double that to 64. See *MaxChecksumBlocks* on page 126.
- **Target write speed**—In high-bandwidth environments, Double-Take throughput is most often limited by the write speed of the target disks. Accordingly, optimizing the target disks for write

performance will often increase Double-Take performance, particularly for full mirrors and high loads of replication. Using RAID 0 and/or RAID 1 instead of RAID 5 on the target disks will improve the target write performance, as well as allocating some (or all) of the I/O controller's cache memory to write operations.

- **TCPBufferSize**—Network throughput is directly related to the TCP buffer size and the network latency of the LAN or WAN connection. By default, Double-Take is configured for a 1Gbit LAN network. If you are replicating across a different LAN network or a WAN network, adjust the TCP buffer size accordingly. For example, for a 100Mbit LAN, the value should be around 37500, and for a WAN, the value should be around 130000. See *TCPBufferSize* on page 148 or the technical support article 31662.
- **Windows MTU**—The Maximum Transmission Unit (MTU) is the largest amount of data, a packet, that can be transferred in one physical frame on a network. If the MTU is too high, you may get fragmented packets which can slow down Double-Take mirroring and replication and can possibly cause lost Double-Take connections. Use the ping command with the -f -l 1500 options. If you receive a response that packets need to be fragmented, you should lower your MTU value. See the Microsoft article 314825 for details on specifying the MTU value.
- **Disable root encryption**—If the top-level folders in your jobs are not encrypted, you can gain a performance improvement by disabling root encryption. See *EnableRootEncryption* on page 116.

General manageability

- **Temporary files**—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, some unexpected data may be replicated if applications use the \Local Settings\Temp directory to store data. This could result in significant amount of unnecessary data replication on large file servers. Additionally, the \Local Settings\Temporary Internet Files or \AppData\Local\Microsoft\Windows\Temporary Internet Files directories 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. You may want to consider excluding temporary data like this, however it is important to know how applications may use these temporary files. For example, Microsoft Word creates a temporary file when a document is opened. When the user closes the file, the temporary file is renamed to the original file and the original file is deleted. In this case, you must replicate that temporary file so that Double-Take can process the rename and delete operations appropriately on the target.
- **E-mail notification**—Enable e-mail notification through the *E-mail notification configuration* on page 97 so that you are notified when a Double-Take message is written to the Event log for that server.
- **Target path blocking**—Target path blocking prevents the modification of the copy of the source data on the target until failover has occurred or protection is disabled. This can be can be configured for some job types or for all jobs to a target through the *Target server properties* on page 95.
- **File difference mirrors**—When performing a file difference mirror, you may not want to use the **Mirror only if the file on the source is newer than the copy on the target** option, unless you know for certain that it is acceptable. This is particularly true for database applications and domain controllers since it is critical that all files, not just some of them that might be newer, are mirrored.
- **Disable attribute replication**—On servers where the file permissions need to be different on the source and target, you can disable the replication of file attributes. When attribute replication is disabled, files on the target can inherit permissions from the parent directory on the target. See *TGDisableAttributeReplication* on page 150.

Anti-virus protection

- **Double-Take queue**—Exclude the Double-Take queue directory on the source and target from any real-time scanning or scheduled system scans. If a queue file is deleted by a process other than Double-Take, unexpected results may occur, including an auto-disconnect due to the loss of queued data. The files in the source queue directory have already been scanned (cleaned, deleted, or quarantined) in their original storage location. The files in the target queue have already been scanned (cleaned, deleted, or quarantined) on the source.
- **Target data**—Exclude the copy of the source data stored on the target from any real-time scanning or scheduled system scans. The files have already been scanned (cleaned, deleted, or quarantined) on the source. 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 being protected. 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. Additionally, if the virus protection software cleans the file, an operation to clean the file will likely also be replicated from the source, which may result in file corruption.

Hardware configurations

- **NIC teaming**—If you are using NIC teaming, set it up for fault tolerance, not load balancing.
- **Device drivers**—Keep your hardware device drivers, especially NIC drivers, up-to-date.
- **Port speed and duplex**—Set static values for port speed and duplex on NICs and switches, if possible.

Full server optimizations

Review the following optimizations when you are using a full server protection job.

- **Third machine to run Double-Take Console**—Ideally, you should use a third machine to run the Double-Take Console and set up protection and to perform failover and reverse. If you do not use a third machine, you may need to remove and reinsert your servers (using reserved IP addresses) into the console. If you use a third machine, it must be able to communicate with the reserved IP addresses.
- **Single gateway**—If you are using multiple IP addresses on Windows 2003, do not use multiple gateways.
- **NIC configuration**—If you are planning to failover the IP address of the source, use a separate NIC and separate network for a Double-Take reserved IP address that will not be failed over. If you are unable to do that and just one NIC is used for both production and reserved IP addresses, disable DNS registration on the NIC. If you are not going to failover the IP address of the source, an additional NIC and address is not necessary. In this case, Double-Take will block the DNS record for that address while it is failed over.
- **Single NIC**—If you have to use only one NIC, disable DNS registration and ensure the reserved IP address is first in the list of IP addresses. For Windows 2003, if the failover IP address is listed first and it cannot come online due to IP address conflicts, subsequent IP addresses may not come online. Also, the original source IP address may need to be removed manually after bringing the original source back online.
- **Disabling DNS registration**—Disabling DNS registration for the reserved IP address ensures that an end-user is not communicating to the original source when it is failed over because two different DNS records will point to two different servers.
- **Unnecessary target components**—Do not install applications, features or language packs on the target that are not essential to that machine. These can slow the failover process and/or cause mirroring issues.
- **Vendor applications**—Disable or remove any vendor applications that you are not using. These applications may cause issues after failover, like application crashes or even server crashes. These applications may not be in standard application directories, like C:\Dell or C:\cqpsystem, and those directories should be added to the list of staged folders. Also, they may use resources (like C:\INetPub), which may be need after failover. Those directories would also have to be staged.
- **Hardware maintenance**—If you replace a NIC after you have established full server protection, you should delete and re-create your job.

Application optimizations

Review the following optimizations when you are using an Exchange Server or SQL Server job or if you are protecting any other application on your source.

- See *General applications* on page 836
- See *Exchange* on page 837
- See *SQL* on page 837

General applications

- **Application services on the target**—Ensure that all application services on the target are stopped and set to manual.
- **Connection mappings**—When protecting an application server, select the **One To One** mapping when creating a files and folders job.
- **File difference mirrors**—When performing a file difference mirror, you may not want to use the **Mirror only if the file on the source is newer than the copy on the target** option, unless you know for certain that it is acceptable. This is particularly true for database applications and domain controllers since it is critical that all files, not just some of them that might be newer, are mirrored.
- **Mirror using write-through mode**—You can set a Double-Take driver option which allows mirroring to open files in write-through mode which is the same way many applications, such as Exchange and SQL, open files. This may improve performance on application servers. Under HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\RepKap\Parameters, add a DWORD value called DisableKfaiCaching, and set it to 1.
- **Database backups**—Pause the target while you perform a backup of database files stored on the target because the database and log files must be backed up when they are at the exact same point in time. For example, say the back up of the file mydatabase.mdf begins on the target. While the backup program has access to the file, Double-Take cannot write to the file. When the backup completes, Double-Take writes to the file. Double-Take also writes to the corresponding mydatabase.ldf file. When the backup gets to the mydatabase.ldf file, it no longer matches the .mdf file. The database may require special startup procedures, which may result in incomplete transactions being left in the database or data loss. To workaround this scenario, pause the target before starting the backup and then resume the target when the backup is complete.

Exchange

- **Orphan log files**—Exchange database and log files should be synchronized between the source and target, including the removal of any orphan log files on the target. Exchange may not recover the databases correctly if orphan log files are present on the target. (Orphan files are files that exist in the copy of the source data on the target, but are not on the source. This usually occurs when a file is deleted on the source when there is no Double-Take connection.) Make sure you configure your Exchange protection to remove orphan files.
- In most cases, the length of time it takes for the Exchange services to start is directly proportional to the number of log files present. The number of log files can be minimized through frequent Exchange backups on the source. When Exchange performs a full backup it will set the archive bit and mark the committed log files for deletion. Exchange should delete the committed log files during the normal Exchange maintenance if the archive bit has been set.

SQL

- **Memory**—Typically, SQL uses all available memory. You may want to consider limiting SQL memory usage to allow the Double-Take service to function without running out of memory.
- **Temp database**—Check with your application vendor to determine if the temp database is used and needed. If it is not needed, you can exclude it from replication. For example, SQL Server re-creates the tempdb database file each time it starts, so any tempdb data that gets replicated to the target will never get used. Writes to the tempdb database may account for a significant percentage of writes to all SQL Server files, so excluding the tempdb files may result in much less replication traffic. If the database application you are using uses the temp database file (for example in Prophecy, PeopleSoft, and BizTalk) or if you are uncertain, do not exclude it from replication.
- **SQL service account**—Configure the source and target to use the same domain account to start the SQL services, if possible. This eliminates the need to move SQL Service Principal Names (SPNs) during failover and failback. If you have to use different accounts, Kerberos authentication will require the Service Principal Names to be failed over. See the technical support article 33398 for details on failing over the Service Principal Names.

Chapter 21 Security

To ensure protection of your data, Double-Take offer multi-level security using native operating system security features. Privileges are granted through membership in user groups defined on each machine. To gain access to a source or target, the user must provide a valid operating system user name and password and the specified user name must be a member of one of the Double-Take security groups. Once a valid user name and password have been provided and the source or target has verified membership in one of the security groups, the user is granted appropriate access to the source or target and the corresponding features are enabled in the client. Access is granted on one of the following three levels.

- **Administrator Access**—All features are available for that machine.
- **Monitor Access**—Servers and statistics can be viewed, but functionality is not available.
- **No Access**—Servers appear in the clients, but no access to view the server details is available.

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

- *Adding users to the security groups* on page 839
- *Changing the account used to run the Double-Take service on Windows servers* on page 840

Adding users to the security groups

The security groups are automatically created during the installation process.

- **Windows servers**—The Double-Take Admin and Double-Take Monitors groups are automatically created and the local administrator and domain administrator are automatically added to the Double-Take Admin group during installation.
- **Linux servers**—The groups can be local or LDAP (Lightweight Directory Access Protocol). The groups are called dtadmin (default group ID 501) and dtmon (default group ID 502). During the installation, the user root is automatically added to the dtadmin group.

If Double-Take is installed on a Windows member server, it will use the local groups. If an Active Directory user is granted access to the Active Directory Double-Take Admin or Double-Take Monitors groups, the user or domain group must also be granted access to the local Double-Take groups. If Double-Take is installed on a Windows domain controller, the Active Directory group will provide sufficient access. The groups are created in the Users Container and need to stay here. If the groups are not there, users will be unable to log into Double-Take on the domain controller.

Users that need administrator access to Double-Take must be added to the Double-Take Admin or dtadmin group, depending on the operating system. Users that need monitor only access must be added to the Double-Take Monitors or dtmon group, depending on the operating system. In both cases, local users, domain users, or global groups may be added to the local groups.

See your Windows documentation for instructions on adding, deleting, or modifying users in a security group.

For Linux servers, use the following steps.

1. Run the DTSetup command from the shell prompt. The command is case-sensitive. Do not run DTSetup using the sudo command. Use a real root shell to launch DTSetup instead, either by logging in as root on the console or by using the login session of a non-privileged user to run su - to start a root shell.
2. Select **Setup tasks**.
3. Select **Add/Remove users to Double-Take groups**.
4. Select the necessary menu options to add or remove groups to the administrator or monitors group as needed, and specify the user name when prompted.
5. When you have completed your security group modifications, press Q as many times as needed to return back to the main menu or to exit DTSetup.

Changing the account used to run the Double-Take service on Windows servers

By default, the Double-Take service on Windows servers is configured to log on as the system account. If you want to select a specific account to run the service, use these instructions.

If you are protecting an entire server, you cannot modify the account used to run the Double-Take service. Otherwise, the full server protection will not function correctly.

1. Modify the user account that the Double-Take service is using.
   a. Open the Double-Take service properties and select the **Log On** tab, select **This Account**, and enter a valid domain account.
   b. Enter the password for this account.
   c. Click **OK** to save these settings.
2. Grant an additional user right to the account you are using to run the Double-Take service.

   

   If domain-level policy settings are defined (through **Domain Security Policy**, **Security Settings**, **Local Policies**, **User Rights Assignment**), they will override local policy settings.

   a. Select **Local Security Policy** from Administrative Tools.
   b. Expand the **Local Policies** folder and highlight the **User Rights Assignment** folder.
   c. Double-click the option **Act as part of operating system** on the right pane of the screen.
   d. Add the user that you selected to run the Double-Take service and click **OK**.
   e. Exit the Local Security Settings dialog box. This user is now configured to run the Double-Take service.
3. Add the domain account to the local administrator group.
   a. Select **Computer Management** from Administrative Tools.
   b. Expand the **Local Users and Groups** folder and highlight the **Groups** folder.
   c. Right-click on the **Administrators** group on the right pane of the screen and select **Add to Group**.
   d. Click **Add**.
   e. Locate the domain account that you are using for the Double-Take service. Select that account and click **OK**.
   f. Click **OK** to close the Administrators Properties dialog box.
   g. The domain account is now added to the local administrator group. Close the Computer Management window.

Index

G

H

Q

R

S

T

U

V

W