Double-Take®

Version 7.1.2

Double-Take Availability for Hyper-V User's Guide

Notices

Double-Take Availability for Hyper-V User's Guide Version 7.1.2, Thursday, December 15, 2016

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

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

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

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

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

Contents

Chapter 1 Double-Take Availability for Hyper-V overview **4**

Chapter 2 Agentless Hyper-V requirements **5**

Chapter 3 Creating an agentless Hyper-V job **13**

Chapter 4 Simulating protection **35**

Chapter 5 Configuring Hyper-V Pro tip integration for failover notification **36**

Chapter 6 Managing and controlling agentless Hyper-V jobs **38**
- Viewing agentless Hyper-V job details 49
- Validating an agentless Hyper-V job 53
- Editing an agentless Hyper-V job 54
- Viewing an agentless Hyper-V job log 55

Chapter 7 Failing over agentless Hyper-V jobs **57**

Chapter 8 Reversing agentless Hyper-V jobs **59**

Chapter 9 Double-Take Console **60**
- Double-Take Console requirements 62
- Console options 63

Chapter 10 Managing servers **67**
- Adding servers 73
- Providing server credentials 75
- Viewing server details 76
- Editing server properties 78
  - General server properties 79
  - Server licensing 80
  - Server setup properties 82
  - Double-Take queue 86
  - Source server properties 90
  - Target server properties 92
  - E-mail notification configuration 94
  - Script credentials 96
  - Log file properties 97
    - Verification log 99
- Viewing server events 101
- Viewing server logs 102
- Managing snapshots 104
  - Snapshot states 105

Chapter 1 Double-Take Availability for Hyper-V overview

Double-Take Availability ensures the availability of critical workloads. Using real-time replication and failover, you can protect multiple virtual machines on a Windows Hyper-V host. Because Double-Take Availability for Hyper-V is agentless, you can protect your virtual machines without having to install Double-Take on each guest. You identify the virtual machines on your production Hyper-V host, known as the source, and replicate them to a backup Hyper-V host, known as the target. The target server, on a local network or at a remote site, stores a replica of each virtual machine. Double-Take monitors any changes to the source and sends the changes to the replica copies stored on the target. By replicating only the file changes rather than copying an entire file, Double-Take allows you to more efficiently use resources.

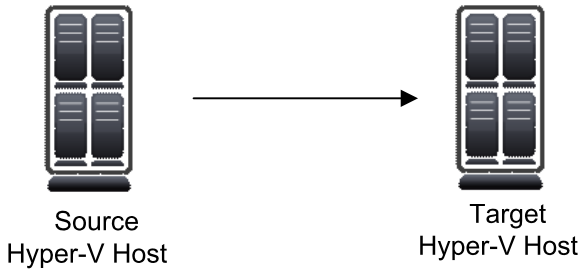

Chapter 2 Agentless Hyper-V requirements

Make sure your servers and virtual machines meet the following requirements.

- **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 that has the Hyper-V role enabled. In addition, you can use Hyper-V Server 2008 R2, Server Core 2008 R2 SP1, 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 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 located in the root of a CSV.
  - 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.
  - The virtual machines' disks cannot reside on storage accessed by a UNC path.
  - 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.
- **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 must have a unique server name.

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

- **Time**—The clock on your Double-Take servers must be within a few minutes of each other, relative to UTC. Large time skews will cause Double-Take errors.
- **Microsoft .NET Framework**—Microsoft .NET Framework version 4.0 Update 3 or later is required on the source and target. (The full .NET 4.0.3 is required, not just the Client Profile.)
- **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.
- **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.
- **Protocols and networking**—Your servers must meet the following protocol and networking requirements.
  - Your servers must have TCP/IP with static IP addressing.
  - IPv4 is the only supported version.
  - If you are using Double-Take over a WAN and do not have DNS name resolution, you will need to add the host names to the local host file on each server running Double-Take.

    - The following requirements apply to clusters.
        - 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.
- **Advanced network support, such as WAN support**—If you want Double-Take to automatically update networking on the guest operating system during failover, for example for WAN support, 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.
- **DNS updates**—Some job types allow you to failover Microsoft DNS records so the source server name resolves to the target IP addresses at failover time. To be able to set up and failover Microsoft DNS records, your environment must meet the following requirements.
    - The source and target servers must be in the same domain.
    - The target must have WMI/DCOM connectivity to any DNS server that you have configured to be updated.
    - Each server's network adapter must have the DNS suffix defined, and the primary DNS suffix must be the same on the source and 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.
    - If you are using a DNS reverse lookup zone , then the forward zone must be Active Directory integrated. Double-Take is unable to determine if this integration exists and therefore cannot warn you during job creation if it doesn't exist. The zone should be set for secure only updates to allow for DNS record locking.

  DNS updates are not supported for Server Core servers or NAT environments.
- **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.
- **Clusters**— Make sure your cluster meets the following 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 on Cluster Nodes are two resources you can start with. There are many other resources available on

the Microsoft TechNet web site.

- **Operating system versions**—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.
- **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.
- **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.
- **Owning nodes**—In a cluster configuration, if you add a possible owning node to the protected network name after a job has started, you must stop and restart the job. If you do not, the records for the new node will not be locked. This could cause problems with DNS records if the source cluster nodes are rebooted or the resources are otherwise cycled on the new owning node.
- **Licensing**—Each node in the cluster must have a valid Double-Take Availability license key.
- **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.
- **Upgrades**—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.

- **Snapshots**—Double-Take uses the Microsoft Volume Shadow Copy service (VSS) for snapshot capabilities. You can take and failover to Double-Take snapshots using an agentless Hyper-V job. To use this functionality, your servers must meet the following requirements.
  - **Snapshot location**—Snapshots are taken at the volume level and stored on the target, so be sure that you have enough space on your target for snapshots.
  - **Double-Take installation location**—In order to enable Double-Take snapshots, Double-Take must be installed on the system drive. If Double-Take is not installed on the system drive, snapshots will be disabled when enabling protection.
  - **Server IP address**—If you have specified an IP address as the source server name, but that IP address is not the server's primary IP address, you will have issues with snapshot functionality. If you need to use snapshots, use the source's primary IP address or its name.

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

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

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

  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.

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

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

- **NAT**—Agentless Hyper-V jobs do not support NAT environments.
- **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 tables identify the supported configurations for an agentless Hyper-V job.

Configuration	Description	Supported	Not Supported
One to one active/standby	You can protect virtual machines on an active host to a standby host. The standby host has no active virtual machines, only the replica virtual machines. Replication will occur in only one direction.	X	
One to one active/active	You can protect virtual machines on an active host to another active host. Each host has active virtual machines in addition to replica virtual machines. Replication will be bi-directional between the hosts.	X	
Many to one	You can protect virtual machines on multiple hosts to one single host. This will consolidate your replica virtual machines to a single host.	X	
One to many	You can protect virtual machines on a single host to multiple target hosts, although individual virtual machines can only be protected once to a single target host. The will allow you to separate your replica virtual machines across multiple target hosts.	X	
Chained	You cannot protect virtual machines on host 1 to host 2 and then protect them again from host 2 to host 3.		X
Single server	You can protect virtual machines within the same host or cluster.	X	
Standalone to standalone	Your hosts can be in a standalone to standalone configuration.	X	
Standalone to cluster	Your hosts can be in a standalone to cluster configuration.	X	
Cluster to standalone	Your hosts can be in a cluster to standalone configuration.	X	
Cluster to cluster	Your hosts can be in a cluster to cluster configuration.	X	
Cluster Shared Volumes (CSV) guest level	You cannot protect a virtual machine on a Cluster Shared Volume from within the guest operating system because agentless Hyper-V jobs are host-level jobs.		X

Configuration	Description	Supported	Not Supported
Cluster Shared Volumes (CSV) host level	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. 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.	X	

Virtual Machine to Virtual Machine Configuration	Description	Supported	Not Supported
One to one active/standby	You can protect a single virtual machine to a single, standby replica virtual machine. The replica virtual machine is not active while replication is occurring. The replica virtual machine will become active after failover.	X	
One to one active/active	You cannot protect a single virtual machine to a single, active replica virtual machine. The replica virtual machine is not active until after failover.		X
Many to one	You cannot protect multiple virtual machines to a single virtual machine. Each virtual machine you protect creates only one replica virtual machine.		X
One to many	You cannot protect a single virtual machine to multiple virtual machines. A virtual machine can only be protected once, creating a single replica virtual machine.		X
Chained	You cannot protect a single virtual machine to a replica virtual machine to another replica virtual machine. The first replica virtual machine is not active, and therefore cannot be protected, until after failover.		X
Single server	You cannot protect a single virtual machine to itself.		X

Chapter 3 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 or Linux 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.

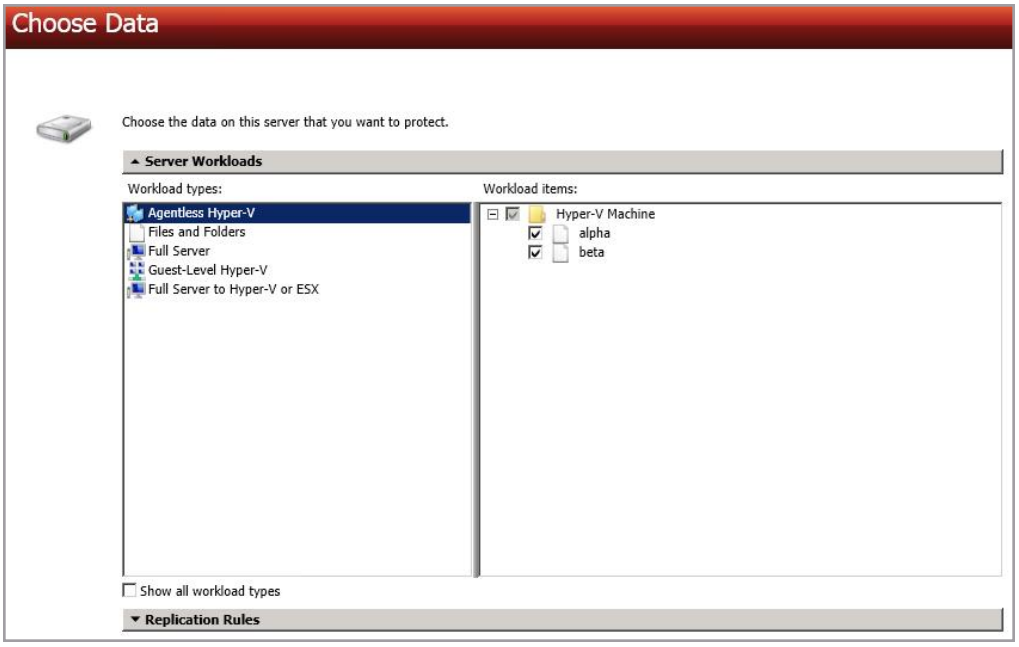

7. Click **Next** to continue.
8. Choose your target server. This is the Hyper-V server or cluster that will store the replicas of the virtual machines from the source.

- **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 33.

- *General* on page 17
- *Replica Virtual Machines* on page 18
- *Replica Virtual Machine Location* on page 19
- *Replica Virtual Machine Configuration* on page 20
- *Replica Virtual Machine Network Settings* on page 21
- *Failover Monitor* on page 22
- *Failover Identity* on page 24
- *Mirror, Verify & Orphaned Files* on page 26
- *Network Route* on page 29
- *Snapshots* on page 30
- *Compression* on page 31
- *Bandwidth* on page 32

General

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

Replica Virtual Machines

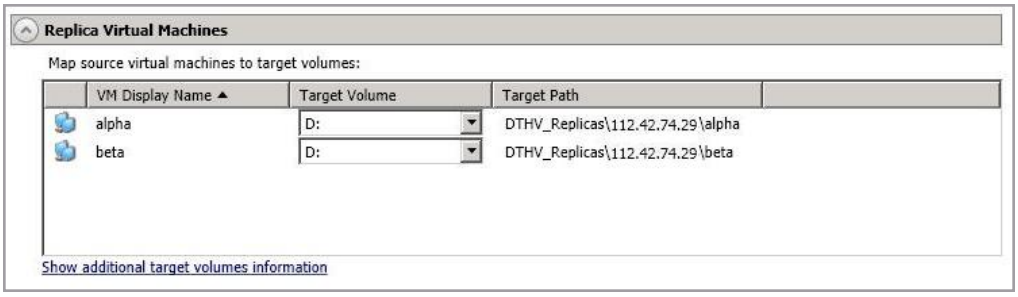

- **Target Volume**—For each virtual machine you are protecting, specify the volume where you want to store the replica virtual machine on the target.
- **Target Path**—For each virtual machine you are protecting, specify a path on the selected **Target Volume** where you want to store the replica virtual 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

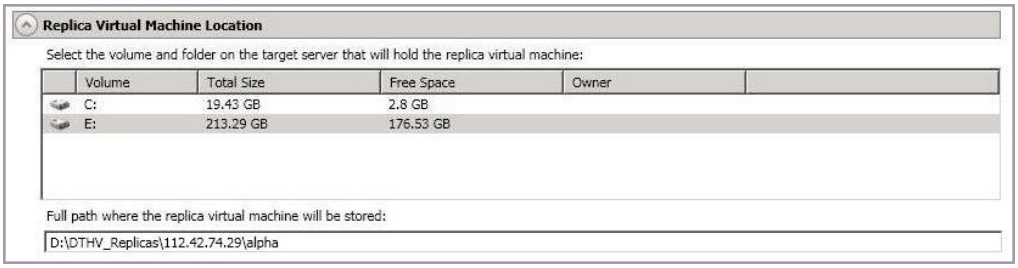

- **Select the volume and folder on the target server that will hold the replica virtual machine**—Select one of the volumes from the list to indicate the volume on the target where you want to store the new virtual server when it is created. The target volume must have enough **Free Space** to store 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 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.

    

    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

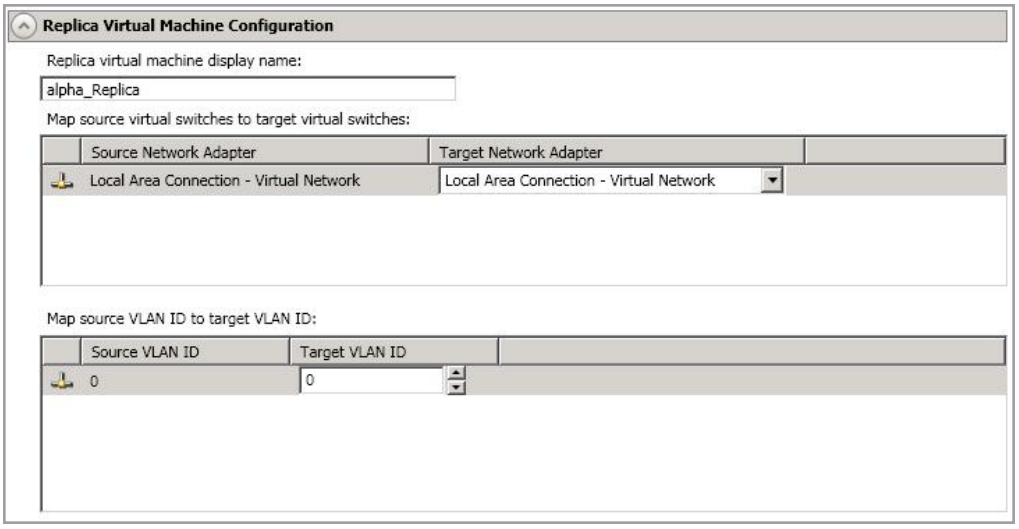

- **Replica virtual machine display name**—Specify the name of the replica virtual machine. This will be the display name of the virtual machine on the host system. You will not see this option if you are protecting more than one virtual machine at a time.
- **Map source virtual switches to target virtual 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. You will not see this option if you are protecting more than one virtual machine at a time.

Replica Virtual Machine Network Settings

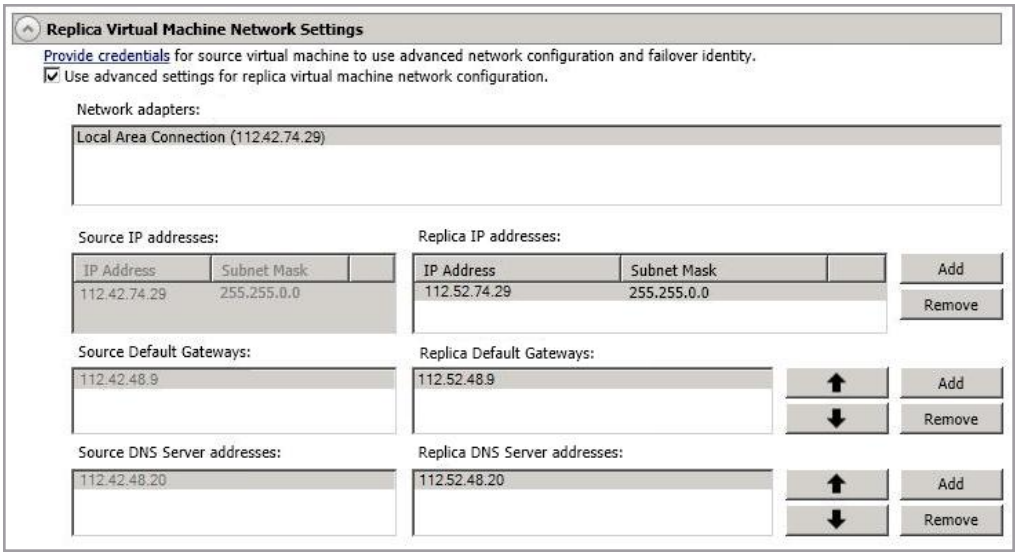

If your virtual machine is powered on and has Integration Services available, this option will allow you to configure advanced settings, which are used primarily for WAN support. Before you can set these options, you must provide credentials for the virtual machine you are protecting. Click the link **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

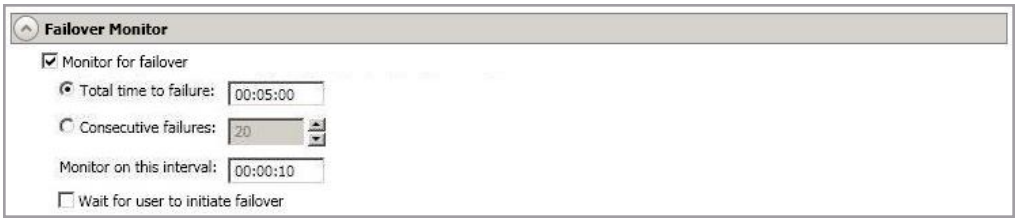

- **Monitor for failover**—Select this option if you want the target to actively monitor the source for a failure. If you disable this option and do not monitor for failover, the target will not actively monitor the source for a failure. In this case, you will have to monitor the source manually on your 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

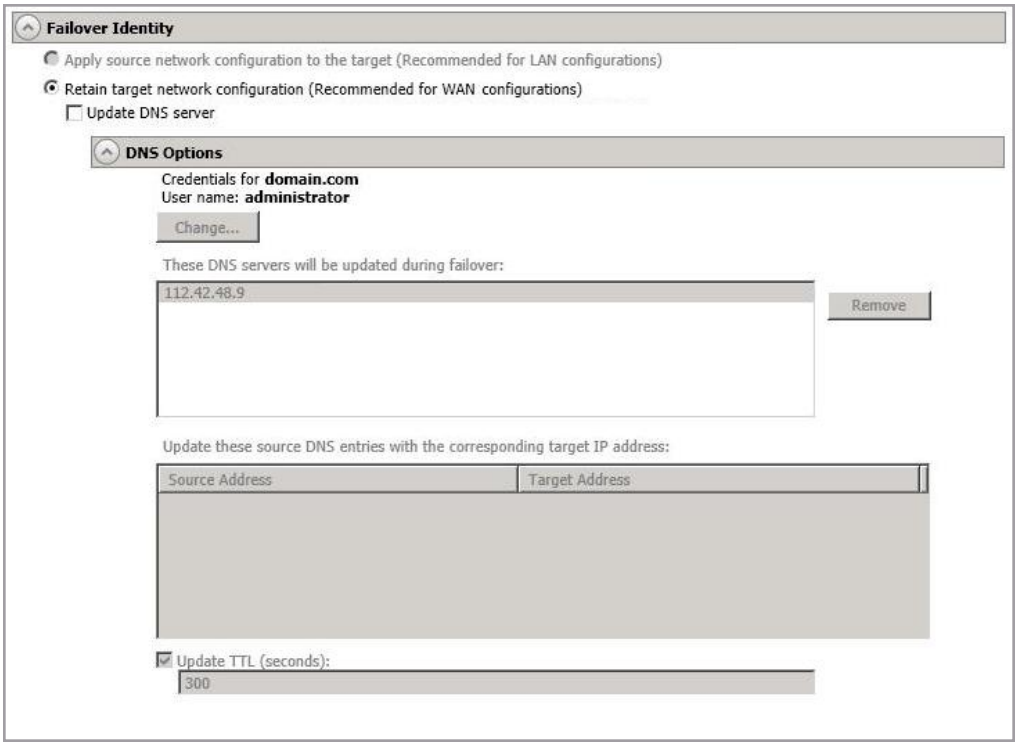

- **Retain target network configuration**—Because the network configuration is set in Replica Virtual Machine Network Settings section, the network configuration is automatically set to retain the target, which is in essence retaining the replica identity. This section is essentially for updating DNS if desired.
  - **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.

    

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

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

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

    The DNS functionality is not available if you have upgraded your job from version 6.0. If you want to use the DNS functionality, you will need to delete your job and create a new one.

    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.

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

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.

Mirror, Verify & Orphaned Files

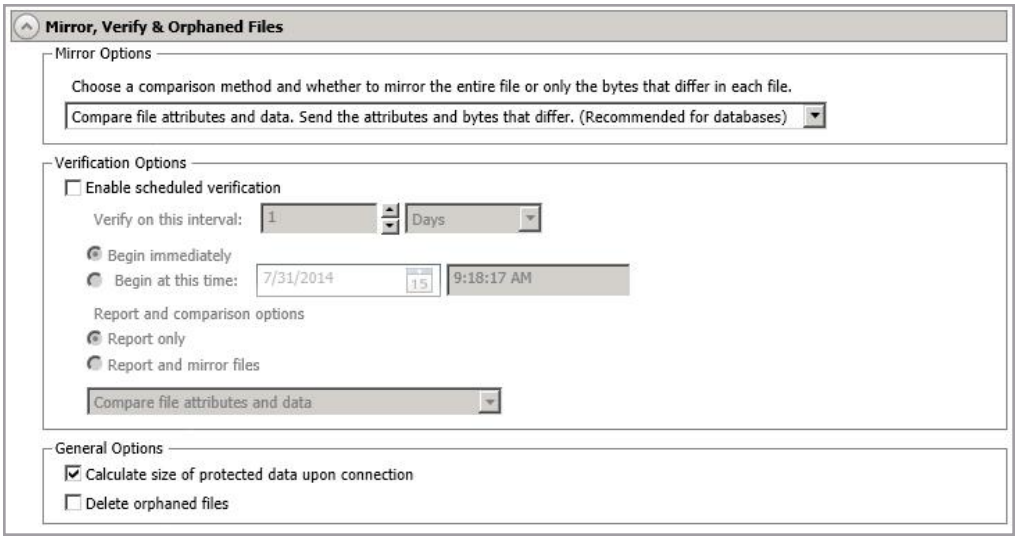

- **Mirror Options**—Choose a comparison method and whether to mirror the entire file or only the bytes that differ in each file.
  - **Do not compare files. Send the entire file.**—Double-Take will not perform any comparisons between the files on the source and target. All files will be mirrored to the target, sending the entire file. This is equivalent to selecting the mirror all files option prior to Double-Take version 7.1.
  - **Compare file attributes. Send the entire file.**—Double-Take will compare file attributes and will mirror those files that have different attributes, sending the entire file. This is similar to selecting the mirror different files and the only if source is newer options prior to Double-Take version 7.1.
  - **Compare file attributes. Send the attributes and bytes that differ.**—Double-Take will compare file attributes and will mirror only the attributes and bytes that are different. This is equivalent to selecting the mirror different files, only if source is newer, and use block checksum options prior to Double-Take version 7.1.
  - **Compare file attributes and data. Send the attributes and bytes that differ.**—Double-Take will compare file attributes and the file data and will mirror only the attributes and bytes that are different. This is equivalent to selecting the mirror different files and use block checksum options prior to Double-Take version 7.1. If you are using a database application on your source, select this option.
- **Verification Options**—Choose if you want to periodically confirm that the source replica data on the target is identical to the actual data on the source. Verification creates a log file detailing what was verified as well as which files are not synchronized. If the data is not the same, you can automatically initiate a remirror, if configured. The remirror ensures data integrity between the source and target.

Because of the way the Windows Cache Manager handles memory, machines that are doing minimal or light processing may have file operations that remain in the cache until additional operations flush them out. This may make Double-Take files

on the target appear as if they are not synchronized. When the Windows Cache Manager releases the operations in the cache on the source and target, the files will be updated on the target.

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

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

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

- **General Options**—Choose your general mirroring options.
  - **Calculate size of protected data upon connection**—Specify if you want Double-Take to determine the mirroring percentage calculation based on the amount of data being protected. If you enable this option, the calculation will begin when mirroring begins. For the initial mirror, the percentage will display after the calculation is complete, adjusting to the amount of the mirror that has completed during the time it took to complete the calculation. Subsequent mirrors will initially use the last calculated size and display an approximate percentage. Once the calculation is complete, the percentage will automatically adjust down or up to indicate the amount that has been completed. Disabling calculation will result in the mirror status not showing the percentage complete or the number of bytes remaining to be mirrored.

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

  - **Delete orphaned files**—An orphaned file is a file that exists in the replica data on

the target, but does not exist in the protected data on the source. This option specifies if orphaned files should be deleted on the target.

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

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 92 for more information.

During a mirror, orphaned file processing success messages will be logged to a separate orphaned file log on the source. This keeps the Double-Take log from being overrun with orphaned file success processing messages. Orphaned files processing statistics and any errors in orphaned file processing will still be logged to the Double-Take log, and during difference mirrors, verifications, and restorations, all orphaned file processing messages are logged to the Double-Take log. The orphaned file log is located in the **Logging folder** specified for the source. See *Log file properties* on page 97 for details on the location of that folder. The orphaned log file is appended to during each orphaned file processing during a mirror, and the log file will be a maximum of 50 MB.

Network Route

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

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.

Snapshots

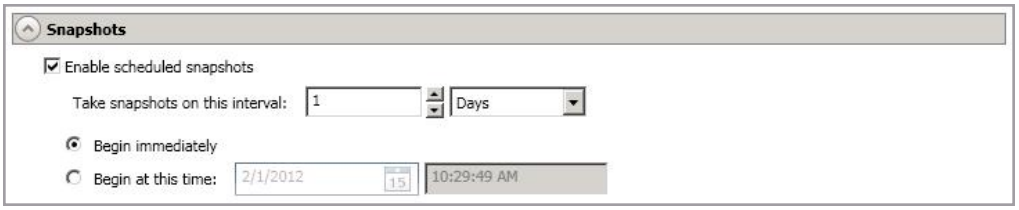

A snapshot is an image of the source replica data on the target taken at a single point in time. You can failover to a snapshot. However, you cannot access the snapshot in VSS to recover specific files or folders.

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

See *Managing snapshots* on page 104 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

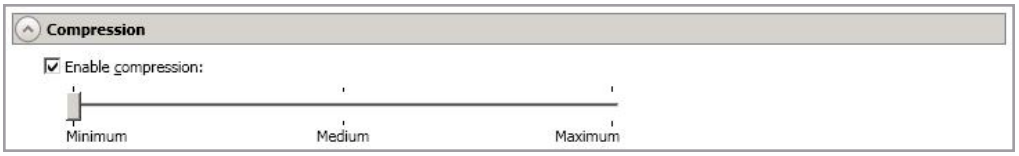

To help reduce the amount of bandwidth needed to transmit Double-Take data, compression allows you to compress data prior to transmitting it across the network. In a WAN environment this provides optimal use of your network resources. If compression is enabled, the data is compressed before it is transmitted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

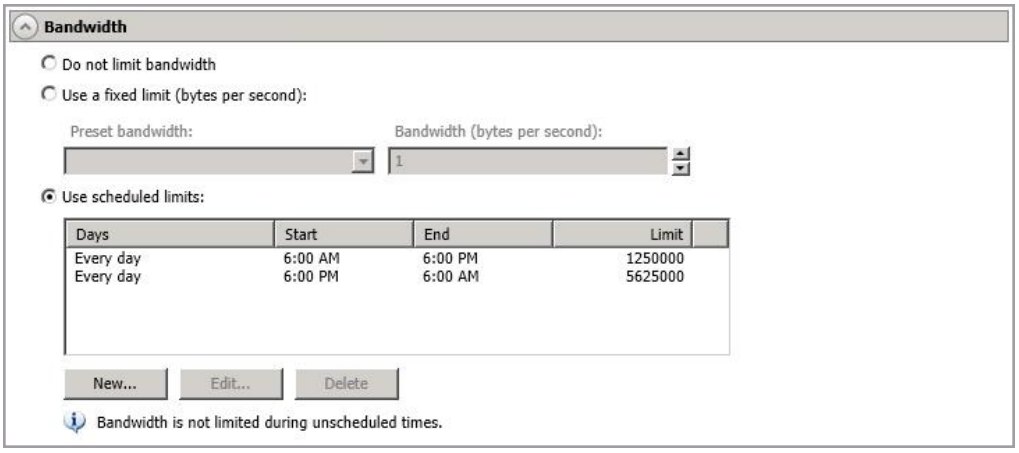

Bandwidth limitations are available to restrict the amount of network bandwidth used for Double-Take data transmissions. When a bandwidth limit is specified, Double-Take never exceeds that allotted amount. The bandwidth not in use by Double-Take is available for all other network traffic.

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 continue. Depending on the error, you may be able to click **Fix** or **Fix All** and let Double-Take correct the problem for you. For those errors that Double-Take cannot correct automatically, you will need to modify the source or target to correct the error, or you can select a different target. You must revalidate the selected servers, by clicking **Recheck**, until the validation check passes without errors.

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

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

13. Once your servers have passed validation and you are ready to establish 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 you are protecting a CSV volume on a Windows 2012 server, you may see event 16400 in the system log during a file rename operation. This error does not indicate a problem with replication or with data integrity on the target and can be ignored.

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 license key), 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.

Chapter 4 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 Double-Take Console to automatically collect this data. It gathers Double-Take log files; Double-Take and system settings; network configuration information such as IP, WINS and DNS addresses; and other data which may be necessary in evaluating Double-Take performance.

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

Chapter 5 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 Double-Take Software 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. *Creating an agentless Hyper-V job* on page 13 to configure failover monitoring and manual intervention.
8. In the event your Hyper-V source fails, you will receive a PRO tip alert.

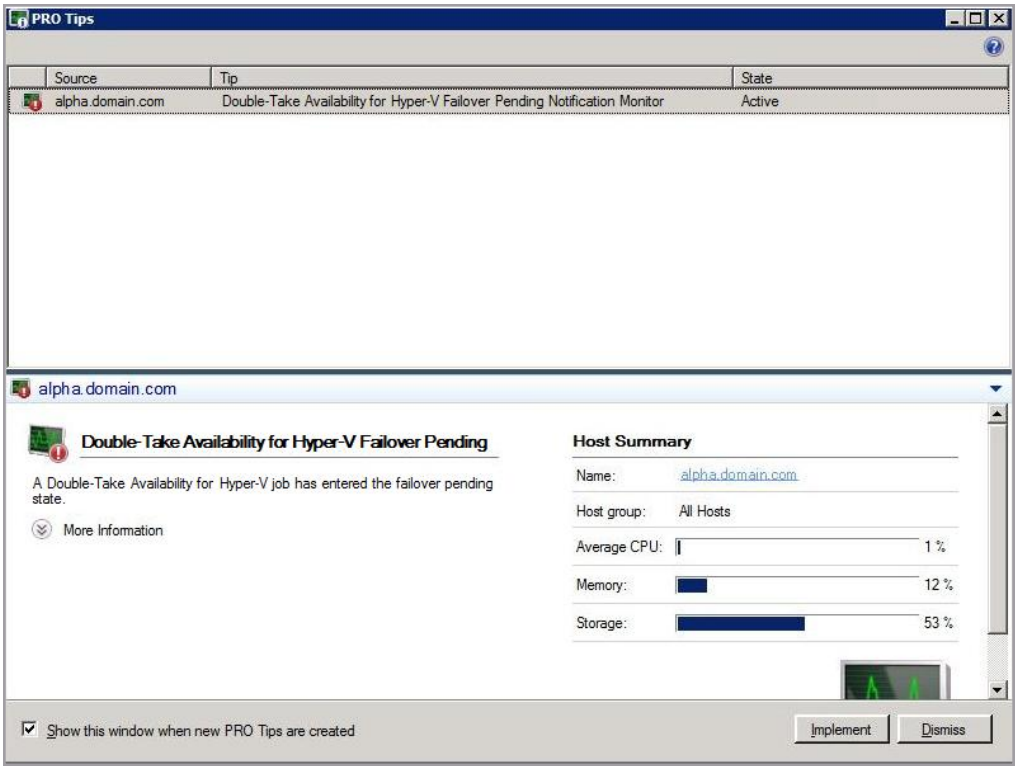

9. You can **Implement** the Pro tip to start failover or you can **Dismiss** it if you do not want to failover.

Chapter 6 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 67), then only the jobs associated with the server or target servers in that server group will be displayed in the right pane.

- *Overview job information displayed in the top pane* on page 38
- *Detailed job information displayed in the bottom pane* on page 41
- *Job controls* on page 43

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 license key), 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.
- **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.
- **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 a license key error.
- **Unknown**—The console cannot determine the status.

Mirror remaining

The total number of mirror bytes that are remaining to be sent from the source to the target. This value may be zero if you have enabled **Mirror only changed files when source reboots** on the *Server setup properties* on page 82.

Mirror skipped

The total number of bytes that have been skipped when performing a difference. These bytes are skipped because the data is not different on the source and target. This value may be zero if you have enabled **Mirror only changed files when source reboots** on the *Server setup properties* on page 82.

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 started.

Recent activity

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

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 75. 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 67.

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 104 for more information.

Failover, Cutover, or Recover 

Starts the failover process. See *Failing over agentless Hyper-V jobs* on page 57 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 59 for the process and details of reversing an agentless Hyper-V job.

Recover 

Recovers the selected DR job. Recovery does not apply to agentless Hyper-V jobs.

Undo Failover or Cutover 

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 57 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.

Other Job Actions 

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

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

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

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

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

  

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

- **Verify**—Even if you have scheduled the verification process, you can run it manually any time a mirror is not in progress.
  - **Report only**—Select this option if you only want to generate a verification report. With this option, no data that is found to be different will be mirrored to the target. Choose how you want the verification to compare the files.
  - **Report and mirror files**—Select this option if you want to generate a verification report and mirror data that is different to the target. Select the

comparison method and type of mirroring you want to use. See the previous mirroring methods described under *Mirror Options*.

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

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

  If you have multiple jobs to the same target, all jobs from the same source will be paused and resumed.
- **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 67.

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 67.

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.
- **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 a license key error.
- **Unknown**—The console cannot determine the status.

Target route

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

Compression

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

Encryption

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

Bandwidth limit

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

Connected since

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

Additional information

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

Mirror status

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

Mirror percent complete

The percentage of the mirror that has been completed

Mirror remaining

The total number of mirror bytes that are remaining to be sent from the source to the target. This value may be zero if you have enabled **Mirror only changed files when source reboots** on the *Server setup properties* on page 82.

Mirror skipped

The total number of bytes that have been skipped when performing a difference. These bytes are skipped because the data is not different on the source and target. This value may be zero if you have enabled **Mirror only changed files when source reboots** on the *Server setup properties* on page 82.

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 continue. Depending on the error, you may be able to click **Fix** or **Fix All** and let Double-Take correct the problem for you. For those errors that Double-Take cannot correct automatically, you will need to modify the source or target to correct the error, or you can select a different target. You must revalidate the selected servers, by clicking **Recheck**, until the validation check passes without errors.

   Validation checks for an existing job are logged to the job log on the target server.
4. Once your servers have passed validation, click **Close**.

Editing 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 13 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 continue. Depending on the error, you may be able to click **Fix** or **Fix All** and let Double-Take correct the problem for you. For those errors that Double-Take cannot correct automatically, you will need to modify the source or target to correct the error, or you can select a different target. You must revalidate the selected servers, by clicking **Recheck**, until the validation check passes without errors.

   

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

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

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.

Because the job log window communicates with the target server, if the console loses communication with the target server after the job log window has already been opened, the job log window will display an error. This includes a target cluster node roll that causes the job log to be hosted by a new cluster node.

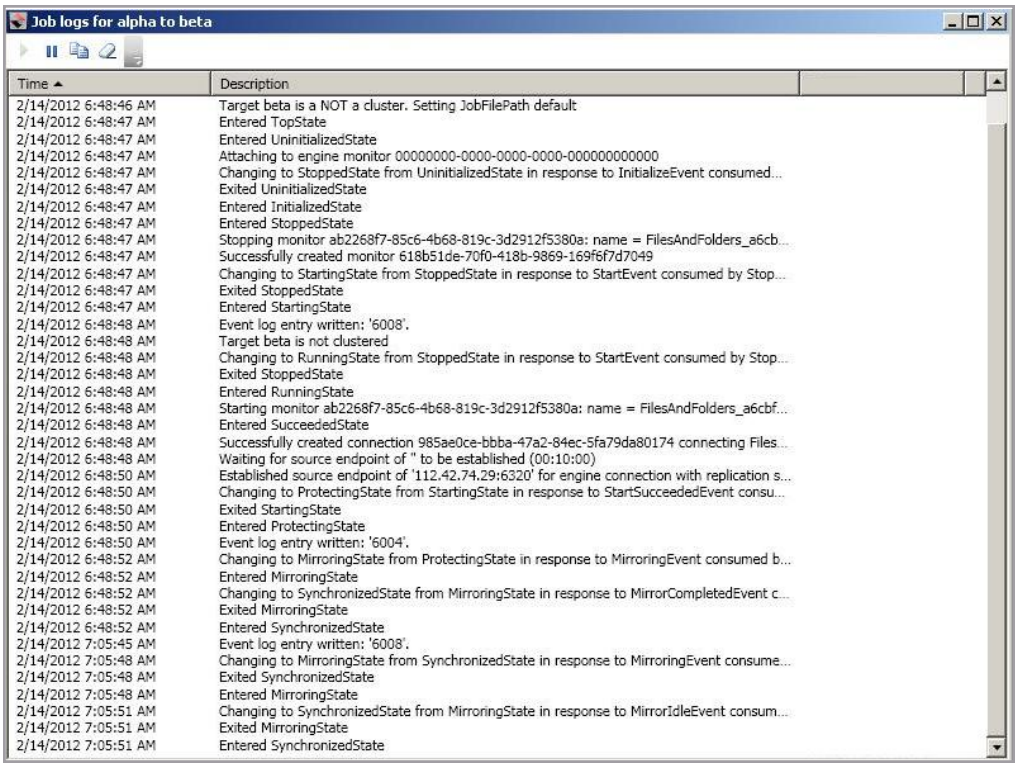

The following table identifies the controls and the table columns in the **Job logs** window.

Start

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

Pause

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

Copy 

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

Clear 

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

Time

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

Description

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

Chapter 7 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, Cutover, or Recover** 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. 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.
    - **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, Cutover, or Recover**.

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 59

Chapter 8 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 9 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 license keys on your servers, and the type of job you are working with.

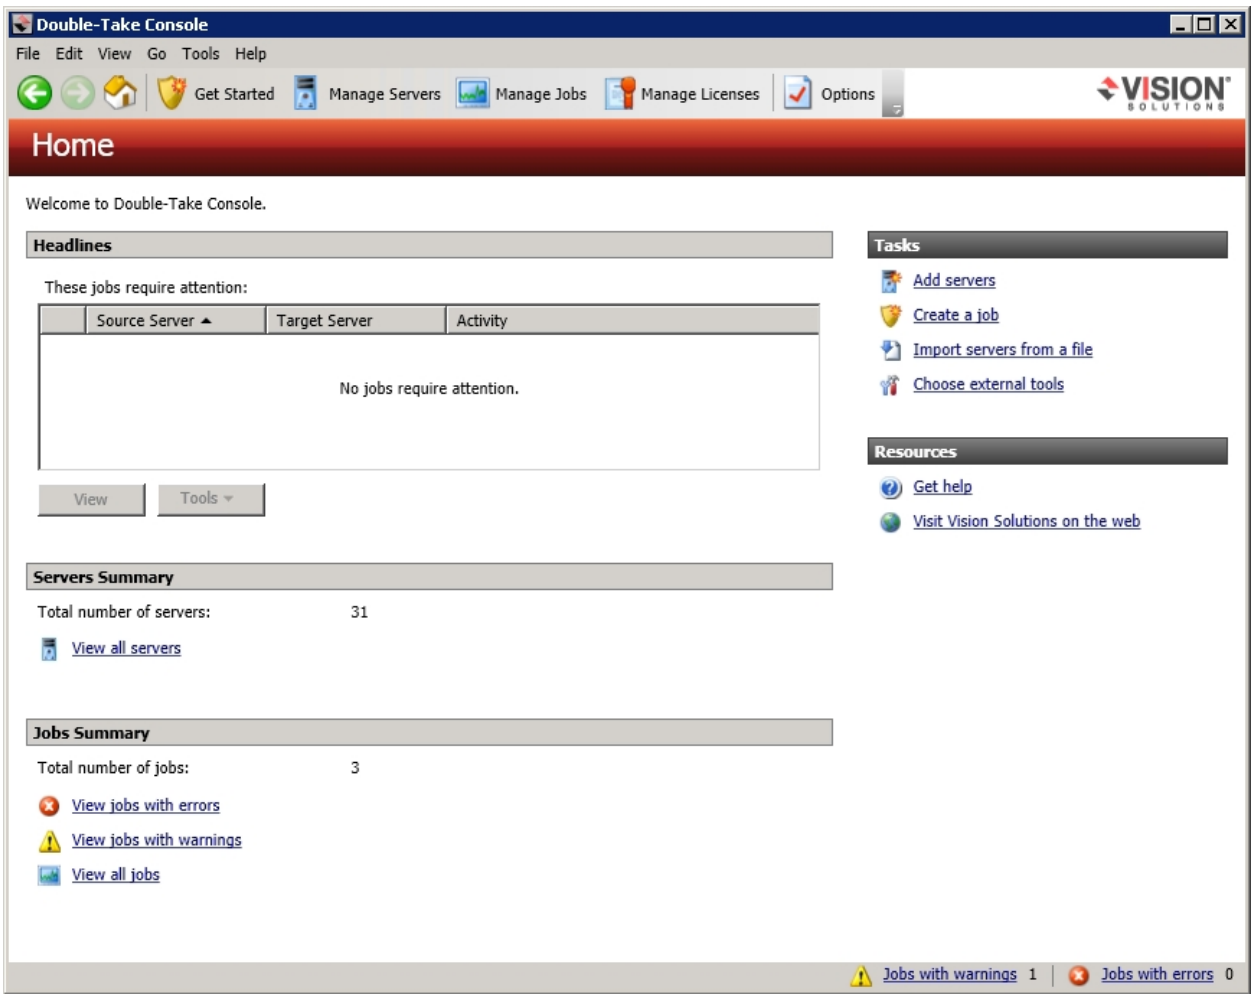

- **Headlines**—The top section gives a quick overview of any jobs that require attention as well as providing quick access buttons.
  - **These jobs require attention**—Any jobs that require attention (those in an error state) are listed. You will see the source and target server names listed, as well 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 Double-Take Software 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 67.
- **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 4.0 Update 3 or later is required. (The full .NET 4.0.3 is required, not just the Client Profile.)
- **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.

If you are running the Double-Take Console on a Windows XP machine and are inserting a Windows 2012 cluster into the console, you must use the IPv4 address to insert the cluster. The console will be unable to connect to the Double-Take Management Service on a Windows 2012 cluster if it is inserted using the name or fully-qualified domain name.

If you are using Windows 2003, you will be unable to open the Double-Take Console context-sensitive help using Internet Explorer. You can access the Double-Take online documentation from the console Help menu using Windows 2003 and Internet Explorer. The help and online documentation are both available in all other Windows versions and in any Windows version using other major browsers.

Console options

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

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

- **Automatically check for updates**—By default, each time the console is started, it will automatically check the Double-Take Software 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 Double-Take Software web site where you can download and install the update.
  - **No update available**—If you are using the most recent console software, that will be indicated. Click **Close**.
  - **No connection available**—If the console cannot contact the update server 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.

- **License Inventory**—This section controls if the console contains a license inventory. This feature may not appear in your console if your service provider has restricted access to it.

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

- **Default Installation Options**—All of the fields under the **Default Installation Options** section are used by the push installation on the **Install** page. The values specified here will be the default options used for the push installation.

  - **Activate online after install completes**—Specify if you want to activate your Double-Take licenses at the end of the installation. The activation requires Internet access from the console machine or the machine you are installing to. Activation will be attempted from the console machine first and if that fails, it wil be attempted from the machine you are installing to. If you choose not to have the installation activate your licenses, you will have to activate them through the console license inventory or the server's properties page.
  - **Location of install folders**—Specify the parent directory location where the installation files are located. The parent directory can be local on your console machine or a UNC path.
    - **Windows**—Specify the parent directory where the Windows installation file is located. The default location is where the Double-Take Console is installed, which is \Program Files\Vision Solutions\Double-Take. The console will automatically use the \i386 subdirectory for 32-bit installations and the \x64 subdirectory for 64-bit installations. These subdirectories are automatically populated with the Windows installation files when you installed the console. If you want to use a different location, you must copy the \i386 or \x64 folder and its installation file to the different parent directory that you specify.

- **Linux**—For Linux servers, you have two choices.
  - If you copied the Linux installation files from your download to your Double-Take Console installation location, you must make sure they are in a \Linux subdirectory under the parent directory you specified for **Location of install folders**. Copy the Linux .deb or .rpm files from your download to the \Linux subdirectory. Make sure you only have a single version of the Linux installation files in that location. The push installation cannot determine which version to install if there are multiple versions in the \Linux subdirectory.
  - If you have already deployed your Linux virtual recovery appliance, specify the UNC path to the installers share on the appliance. For example, if your appliance is called DTAppliance, use the path \\DTAppliance\installers for the **Location of install folders**. The console will automatically use the installation files in the \Linux subdirectory of this share location.

- **Default Windows Installation Options**—All of the fields under the **Default Installation Options** section are used by the push installation on the **Install** page. The values specified here will be the default options used for the push installation.
  - **Temporary folder for installation package**—Specify a temporary location on the server where you are installing Double-Take where the installation files will be copied and run.
  - **Installation folder**—Specify the location where you want to install Double-Take on each server. This field is not used if you are upgrading an existing version of Double-Take. In that case, the existing installation folder will be used.
  - **Queue folder**—Specify the location where you want to store the Double-Take disk queue on each server.
  - **Amount of system memory to use**—Specify the maximum amount of memory, in MB, that can be used for Double-Take processing.
  - **Minimum free disk space**—This is the minimum amount of disk space in the specified **Queue folder** that must be available at all times. This amount should be less than the amount of physical disk space minus the disk size specified for **Limit disk space for queue**.
  - **Do not use disk queue**—This option will disable disk queuing. When system memory has been exhausted, Double-Take will automatically begin the auto-disconnect process.
  - **Unlimited disk queue**—Double-Take will use an unlimited amount of disk space in the specified **Queue folder** for disk queuing, which will allow the queue usage to automatically expand whenever the available disk space expands. When the available disk space has been used, Double-Take will automatically begin the auto-disconnect process.
  - **Limit disk space for queue**—This option will allow you to specify a fixed amount of disk space, in MB, in the specified **Queue folder** that can be used for Double-Take disk queuing. When the disk space limit is reached, Double-Take will automatically begin the auto-disconnect process.

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

- **Default Linux Installation Options**—All of the fields under the **Default Installation Options** section are used by the push installation on the **Install** page. The values specified here will be the

default options used for the push installation.

- **Temporary folder for installation package**—Specify a temporary location on the server where you are installing Double-Take where the installation files will be copied and run.

Chapter 10 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 76.

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.

License Key

The license keys associated with the products licensed for the server. If your license is not valid for the operating system on your server, the license key will be identified as Invalid License Key. There will be no license key listed for those servers that are not licensed, like a VMware server.

Serial Number

The serial number associated with the license key

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. Some options are only available for a single selected server and others are available for multiple selected servers.

Add Servers 

Adds a new server. This button leaves the **Manage Servers** page and opens the **Add Servers** page. See *Adding servers* on page 73.

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 76.

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 75. 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 *Managing and controlling agentless Hyper-V jobs* on page 38.

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

Manage Group Assignments 

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

Install 

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

Uninstall 

Uninstalls Double-Take on the selected server.

Copy 

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

Paste 

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

View Server Events 

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

View Server Logs 

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

Activate Online 

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

already been activated.

Gather Support Diagnostics 

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

View Replication Service Details 

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

Refresh 

Refreshes the status of the selected servers.

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.

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.
   - **Domain**—If you are working in a domain environment, specify the **Domain**.
4. After you have specified the server or appliance information, click **Add**.
5. Repeat steps 3 and 4 for any other servers or appliances you want to add.
6. If you need to remove servers or appliances from the list of **Servers to be added**, highlight a server and click **Remove**. You can also remove all of them with the **Remove All** button.
7. When your list of **Servers to be added** is complete, click **OK**.

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

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.

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

Viewing server details

Highlight a server on the **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.

- **Engine Role**—Source or target server
- **Image Repository Role**—A target for a DR protection job or a source for a DR recovery job
- **Controller Role**—Controller appliance for an agentless vSphere job
- **Replication Appliance Role**—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

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 79—Identifies the server
- *Server licensing* on page 80—Views, adds, and removes license keys
- *Server setup properties* on page 82—Indicates how the server will act on startup and shutdown
- *Double-Take queue* on page 86—Configures the Double-Take queues
- *Source server properties* on page 90—Configures the source server
- *Target server properties* on page 92—Configures the target server
- *E-mail notification configuration* on page 94—Configures e-mail notification
- *Script credentials* on page 96—Specifies credentials to be used when executing custom scripts during mirroring or failover
- *Log file properties* on page 97—Configures log files

Server properties cannot be edited on a cluster.

General server properties

The general server properties identify the server.

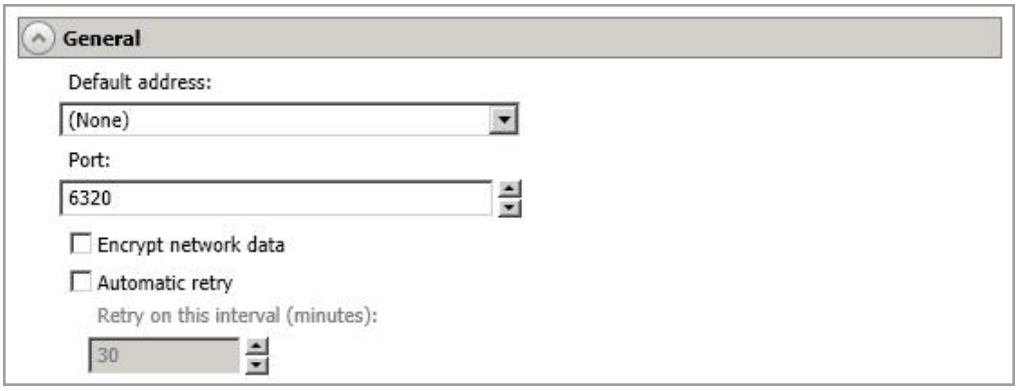

- **Default address**—On a server with multiple NICs, you can specify which address Double-Take traffic will use. It can also be used on servers with multiple IP addresses on a single NIC. If you change this setting, you must restart the Double-Take service for this change to take effect.
- **Port**—The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

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

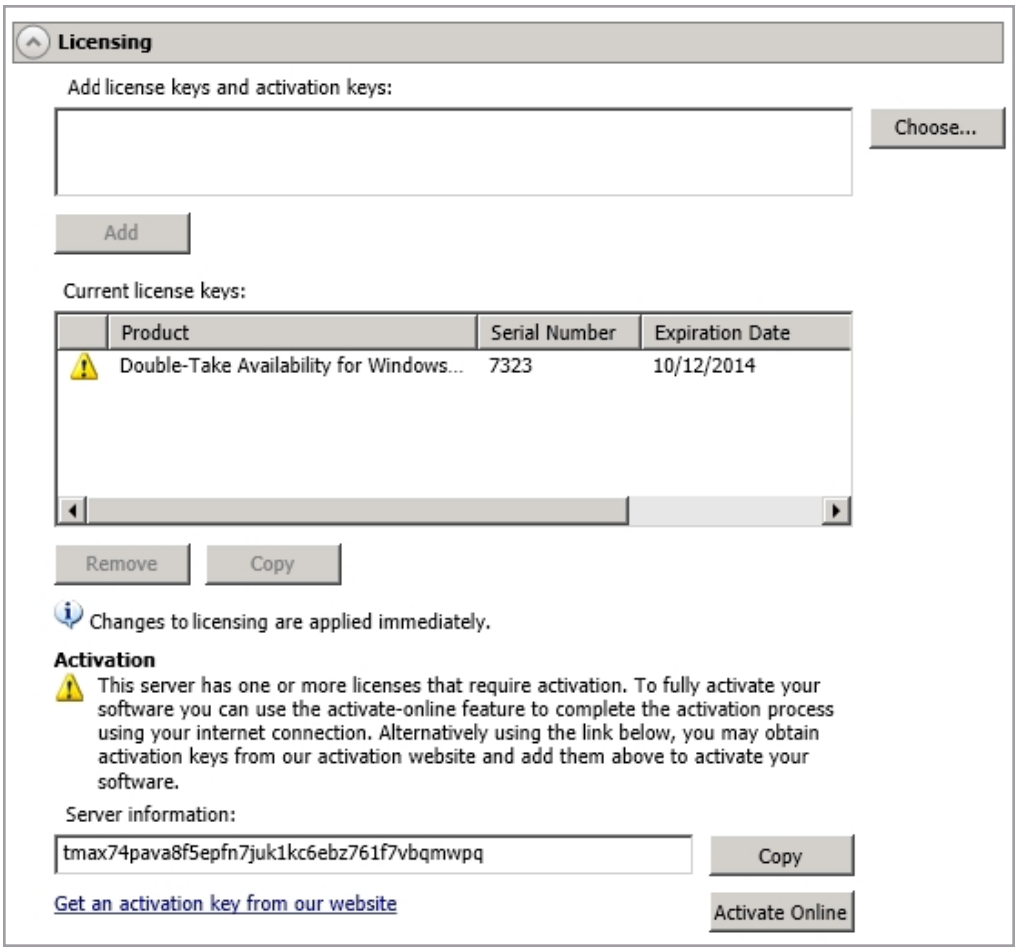

- **Add license keys and activation keys**—Your license key or activation key is a 24 character, alpha-numeric key. You can change your license key without reinstalling, if your license changes. To add a license key or activation key, type in the key or click **Choose from inventory** and select a key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

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

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

If you are replacing an existing license key that has already been activated, you must remove both the old license key and the old activation key. Then you can add a new license key and activate it successfully. If you are updating an existing license key, do not remove the old license key or old activation key. Add the new license key on top of the existing license key.

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

    

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

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

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

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

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

Server setup properties

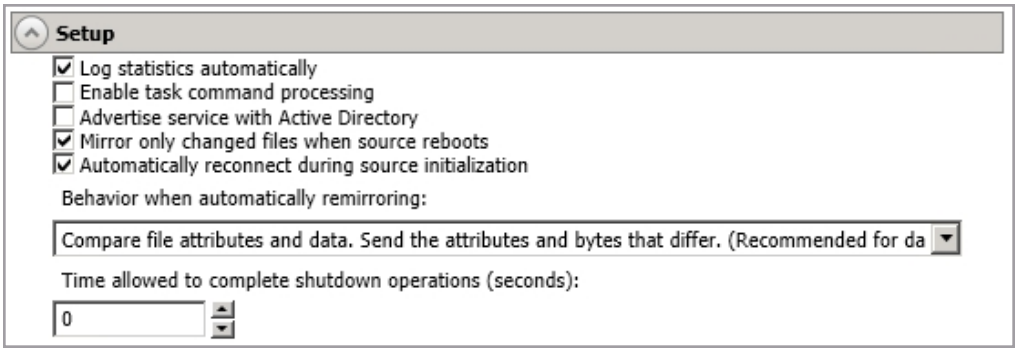

Server setup properties indicate how the server will act on startup and shutdown.

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

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

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

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

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

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

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

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

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

This is equivalent to selecting the mirror different files, only if source is newer, and use block checksum options prior to Double-Take version 7.1. This option is not available for the following jobs.

- Full server
- Full server for Linux
- SQL
- Exchange
- Full server to ESX
- Full server to ESX appliance
- Full server to Hyper-V
- V to ESX
- V to Hyper-V
- Full server migration
- Full server to ESX migration
- Full server to Hyper-V migration
- Full server to Windows DR protection
- Full server to Windows DR recovery

- **Compare file attributes and data. Send the attributes and bytes that differ.**—Double-Take will compare file attributes and the file data and will mirror only the attributes and bytes that are different. This is equivalent to selecting the mirror different files and use block checksum options prior to Double-Take version 7.1. If you are using a database application on your source, select this option.
- **Do not mirror**—Double-Take will not perform any comparisons between the files on the source and target. No files will be mirrored to the target.

- **Time allowed to complete shutdown operations**—This setting indicates the amount of time, in seconds, for the Double-Take service to wait prior to completing a shutdown so that Double-Take can persist data on the target in an attempt to avoid a remirror when the target comes back online. A timeout of zero (0) indicates waiting indefinitely and any other number indicates the number of seconds. The timeout setting only controls the service shutdown caused by Double-Take. It does not control the service shutdown through a reboot or from the Service Control Manager.

Double-Take queue

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

Double-Take Queuing Diagram

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

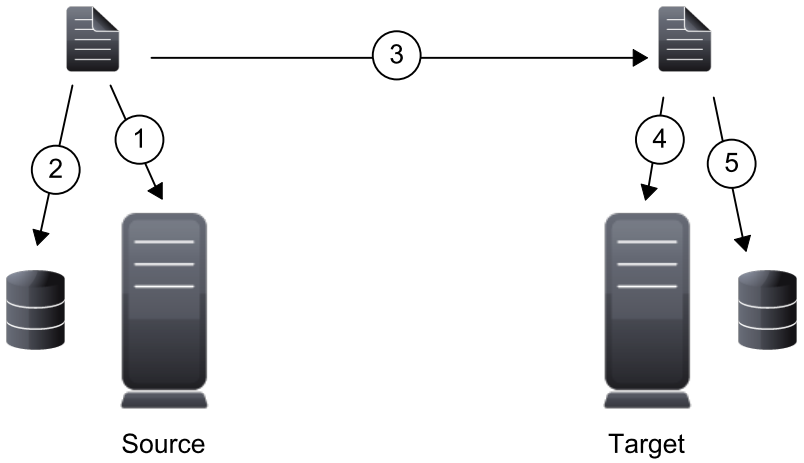

1. If data cannot immediately be transmitted to the target, it is stored in system memory. You can configure how much system memory you want Double-Take to use for all of its processing.
2. When the allocated amount of system memory is full, new changed data bypasses the full system memory and is queued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

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

   ---

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

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

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

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

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

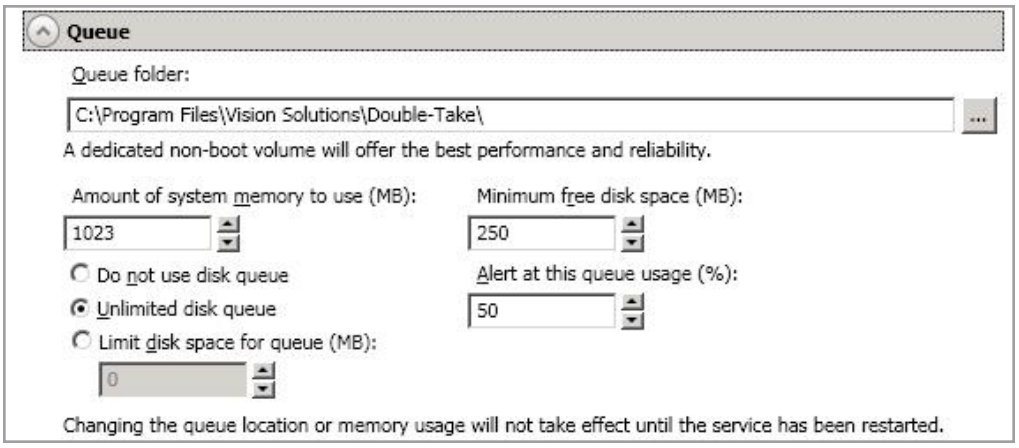

- **Queue folder**—This is the location where the disk queue will be stored. Any changes made to the queue location will not take effect until the Double-Take service has been restarted on the server.

  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.

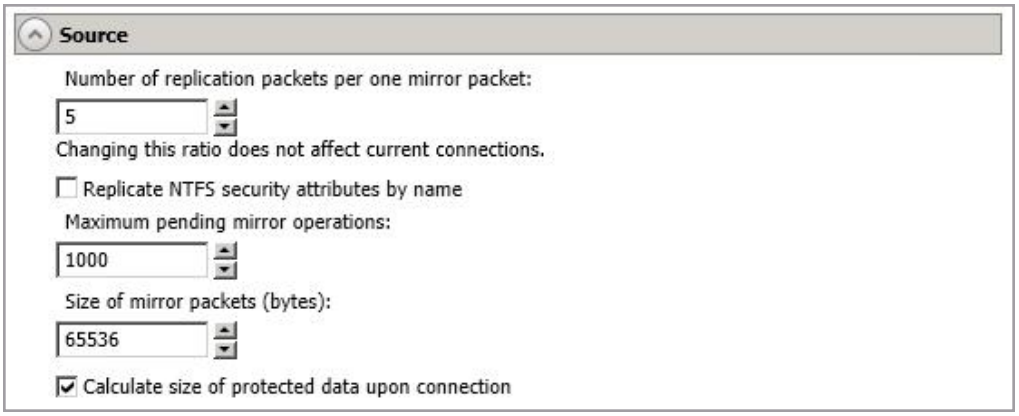

- **Number of replication packets per one mirror packet**—You can specify the ratio of replication packets to mirror packets that are placed in the source queue. The default value (5) allows Double-Take to dynamically change the ratio as needed based on the amount of replication data in queue. If you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

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

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

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

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

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

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

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

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

The calculation time is when the source generates a hard link report that the target will use to process hard links on the target. If you disable the calculation and thus the report generation, the hard link processing that occurs on the target will be ineffective and you may have problems after failover, especially if your source is Windows 2008 or 2012.

Target server properties

These properties are specific to the target server role.

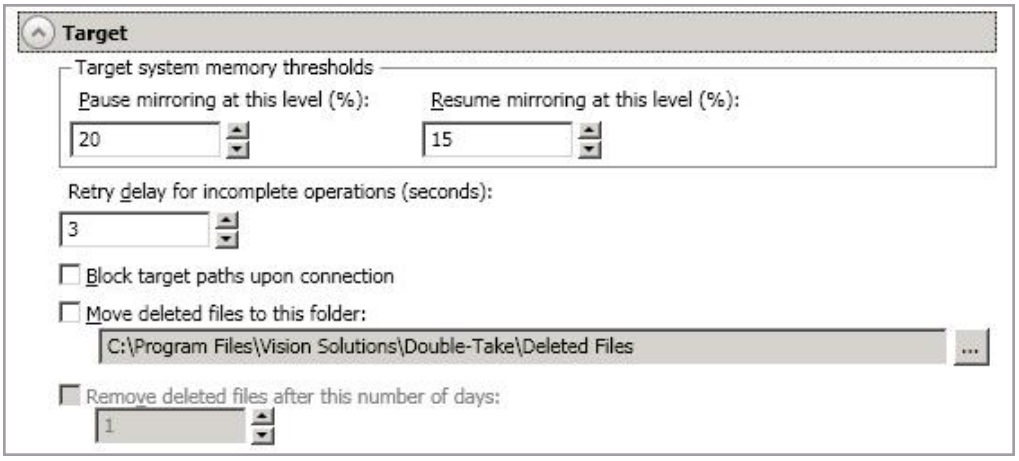

- **Pause mirroring at this level**—You can specify the maximum percentage of Windows system memory that can contain mirror data before the target signals the source to pause the sending of mirror operations. The default setting is 20.
- **Resume mirroring at this level**—You can specify the minimum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

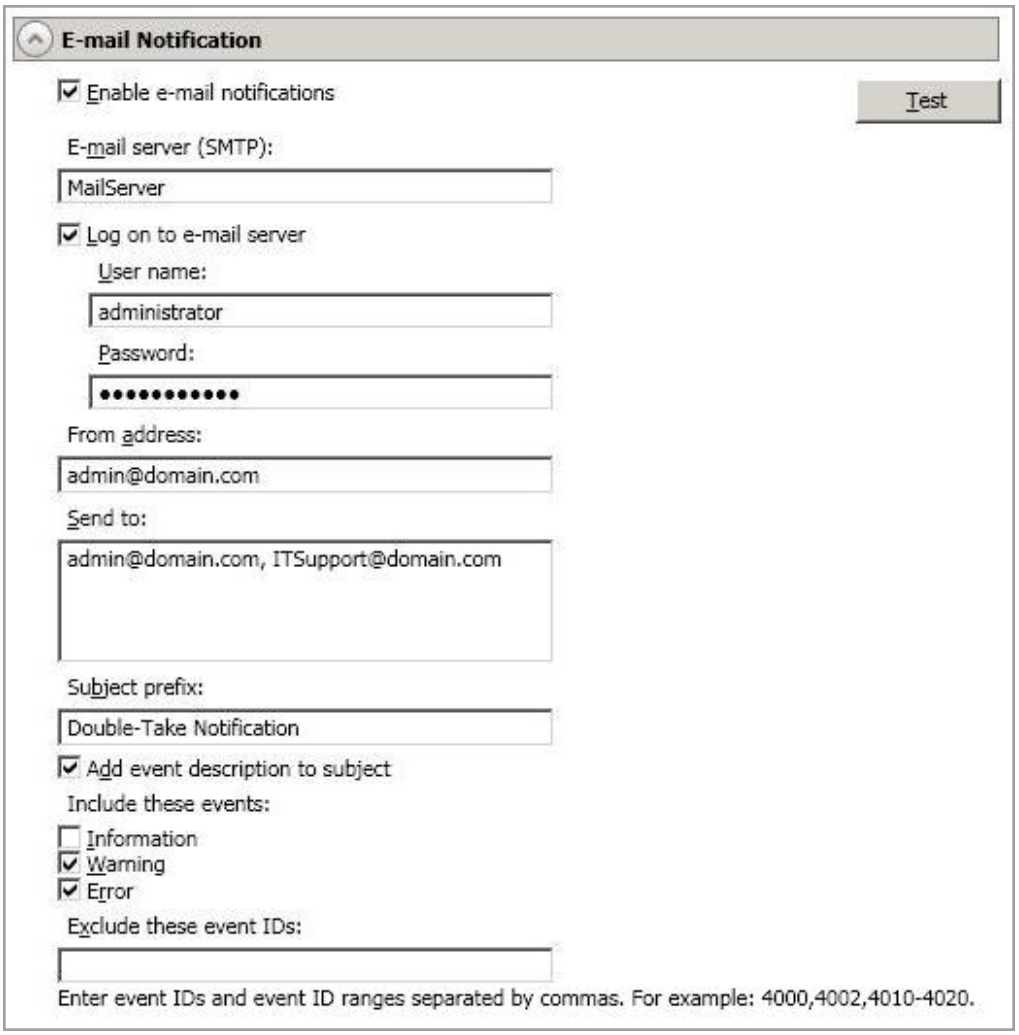

- **Enable e-mail notification**—This option enables the e-mail notification feature. Any specified notification settings will be retained if this option is disabled.
- **E-mail server**—Specify the name of your SMTP mail server.
- **Log on to e-mail server**—If your SMTP server requires authentication,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

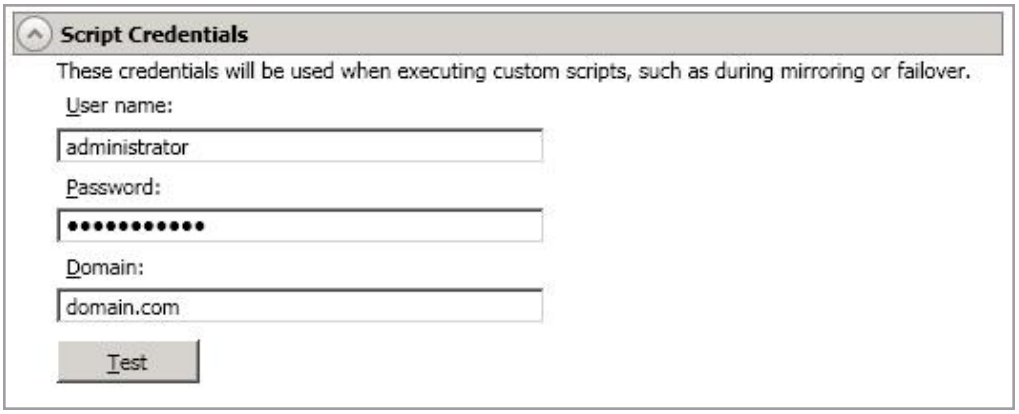

Specify a **User name**, **Password**, and **Domain** to use when running the scripts. If you do not specify any security credentials, the account running the Double-Take service will be used. After you have specified credentials, you can click **Test** to confirm the credentials can be used for a successful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

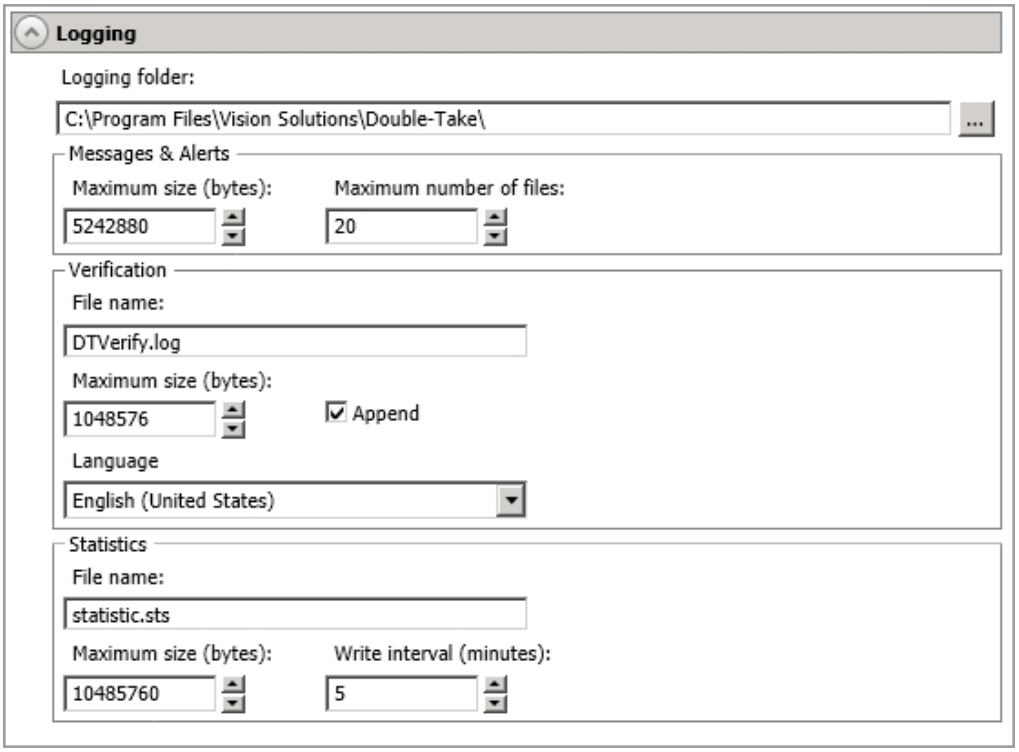

- **Logging folder**—Specify the directory where each of the log files in this section are stored. The default location is the directory where the Double-Take program files are installed.
- **Messages & Alerts**—These settings apply to the service log file.
  - **Maximum size**—Specify the maximum size, 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 20, and the maximum is 999. Once the maximum has been reached, the oldest file will be overwritten.
- **Verification**—The verification log is created during the verification process and details which files were verified as well as the files that are synchronized. See *Verification log* on page 99.
  - **File name**—This field contains the base log file name for the verification process. The job type and a unique identifier will be prefixed to the base log file name. For example, since the default is DTVerify.log, the verification log for a files and folders job will be Files and Folders_123456abcdef DTVerify.log.
  - **Maximum size**—Specify the maximum size, in bytes, of the verification log file. The default is 1048576 bytes (1 MB).
  - **Append**—Enable the **Append** check box if you want to append each verification process to the same log file. If this check box is disabled, each verification process that is logged will overwrite the previous log file. By default, this option is enabled.
  - **Language**—Select the language for your verification log file.

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

Verification log

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

Sample verification log

```
--- VERIFICATION OF CONNECTION 2, CHECKSUM ENABLED (Sales data for alpha --> 206.31.65.40 : 1100) ---
Start Time: 1/24/2016 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 = 1/17/2016 8:21:36 PM Size = 1272 Mask = [0x20]
     Target Attributes: Timestamp = 1/17/2016 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 = 1/17/2016 8:21:37 PM Size = 1272 Mask = [0x20]
     Target Attributes: Timestamp = 1/17/2016 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 = 1/12/2016 3:28:20 PM Size = 17 Mask = [0x20]
     Target Attributes: Timestamp = 1/20/2016 5:05:26 PM Size = 2 Mask = [0x20]
      17 BYTES OUT OF SYNC
Completion Time: 1/24/2016 12:37:44 PM for connection 2 (Sales data for alpha -->
206.31.65.40 : 1100)
Elapsed Time (seconds): 1320.256470
Total Directories Compared: 657
Total Directories Missing: 0
Total Directories Remirrored: 0
Total Files Compared: 120978
Total Files Missing: 0
Total Files Different: 3
Total Files Encrypted: 0
Total Files Remirrored: 1
Total Bytes Skipped: 0
Total Bytes Compared: 18527203678
Total Bytes Missing: 0
Total Bytes Different: 17
Total Bytes Remirrored: 17
Related links and directory attributes have been adjusted.
----- END OF VERIFICATION -----
```

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

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

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

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

Viewing server events

Highlight a server on the **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 the *Reference Guide* for a complete list of Windows event messages.

Viewing server logs

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

- On the **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.

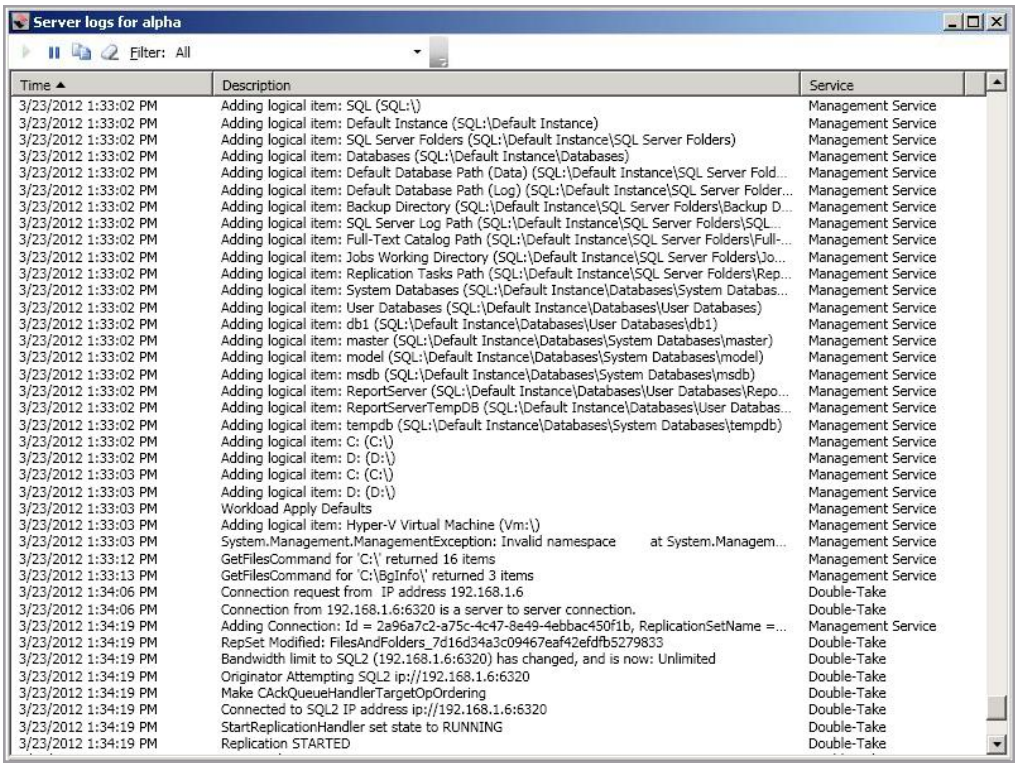

Start

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

Pause

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

Copy 

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

Clear 

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

Filter

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

Time

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

Description

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

Service

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

Managing snapshots

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.
   - **Manual**—This snapshot was taken manually by a user.
3. Click **Take Snapshot** to create a new snapshot for the job.
4. If there is a snapshot that you no longer need, highlight it in the list and click **Delete**.
5. When you have completed your snapshot management, click **Close**.

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

Snapshot states

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

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

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

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

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

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