


Carbonite Migrate for Windows

User's Guide



Notices

Carbonite Migrate for Windows User's Guide Version 8.2.1, Wednesday, October 17, 2018

If you need technical assistance, you can contact CustomerCare. All basic configurations outlined in the online documentation will be supported through CustomerCare. 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 Carbonite Availability and Carbonite Migrate Linux servers. These documents are bound by the same Carbonite license agreement as the software installation.

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

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

© 2018 Carbonite, Inc. All rights reserved.

Contents

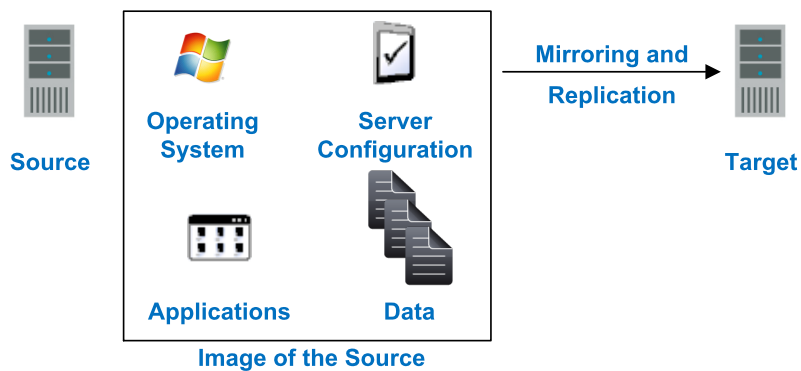
Chapter 1 Carbonite Migrate overview	5
Chapter 2 Requirements	6
Mirroring and replication capabilities	7
Chapter 3 Carbonite Replication Console	12
Carbonite Replication Console requirements	14
Console options	15
Chapter 4 Managing servers	18
Adding servers	28
Providing server credentials	31
Viewing server details	32
Editing server properties	34
General server properties	35
Server licensing	36
Server setup properties	38
Carbonite Migrate queue	41
Source server properties	45
Target server properties	47
E-mail notification configuration	49
Script credentials	51
Log file properties	52
Verification log	54
Viewing server events	56
Viewing server logs	57
Managing VMware servers	59
Chapter 5 Files and folders migration	60
Files and folders migration requirements	61
Creating a files and folders migration job	65
Managing and controlling files and folders migration jobs	83
Viewing files and folders migration job details	93
Validating a files and folders migration job	97
Editing a files and folders migration job	98
Viewing a files and folders migration job log	100
Cutting over files and folders migration jobs	102
Chapter 6 Full server migration	104
Full server migration requirements	105
Creating a full server migration job	112
Managing and controlling full server migration jobs	131
Viewing full server migration job details	141
Validating a full server migration job	145
Editing a full server migration job	146
Viewing a full server migration job log	148
Cutting over full server migration jobs	150
Chapter 7 Full server to ESX migration	153
Full server to ESX migration requirements	154

Creating a full server to ESX migration job	159
Managing and controlling full server to ESX migration jobs	181
Viewing full server to ESX migration job details	191
Validating a full server to ESX migration job	195
Editing a full server to ESX migration job	196
Viewing a full server to ESX migration job log	198
Cutting over full server to ESX migration jobs	200
Chapter 8 Full server to Hyper-V migration	203
Full server to Hyper-V migration requirements	204
Creating a full server to Hyper-V migration job	209
Managing and controlling full server to Hyper-V migration jobs	231
Viewing full server to Hyper-V migration job details	241
Validating a full server to Hyper-V migration job	245
Editing a full server to Hyper-V migration job	246
Viewing a full server to Hyper-V migration job log	248
Cutting over full server to Hyper-V migration jobs	250
Chapter 9 Simulating migration	253
Chapter 10 Special network configurations	254
Firewalls	255
IP and port forwarding	256
Macintosh shares	259
NFS Shares	260

Chapter 1 Carbonite Migrate overview

Carbonite Migrate is a comprehensive migration solution. It allows you to move an entire server, known as a source, by mirroring an image of that source to another server, known as the target. The source and target servers can be physical or virtual. The image of the source contains the server's system state (the server's configured operating system and applications) and all of the source server's data. You can also migrate just a source's data, in which case the target's system state (the target's configured operating system and applications) will be used with the source's data.

Carbonite Migrate uses patented data replication technology that allows users to continue accessing and changing data during the migration. As changes are made on the source, replication keeps the image of the source stored on the target up-to-date. Carbonite Migrate replicates, in real-time, only the file changes, not the entire file, allowing you to more efficiently use resources. When you are ready to cutover to the new server, Carbonite Migrate applies the source system state and after a reboot, the source is available and running on what was the target server hardware.



Chapter 2 Requirements

Each Windows server must meet specific requirements depending on the job type you will be using. See the requirements section for each of the job types for those specific requirements.

- *Files and folders migration requirements* on page 61
- *Full server migration requirements* on page 105
- *Full server to ESX migration requirements* on page 154
- *Full server to Hyper-V migration requirements* on page 204

Mirroring and replication capabilities

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

Carbonite Migrate does not mirror or replicate items that are not stored on the file system, such as physical volume data and registry based data. Additionally, Carbonite Migrate does not mirror or replicate NTFS extended attributes, registry hive files, Windows or any system or driver pagefile, system metadata files (\$LogFile, \$Mft, \$BitMap, \$Extend\\$\UsnJrnl, \$Extend\\$\Quota, and \$Extend\\$\ObjId), or the Carbonite Migrate disk-based queue logs. The only exception to these exclusions is for the full server job types. If you are protecting your system state and data using full server protection, Carbonite Migrate will automatically gather and replicate all necessary system state data, including files for the operating system and applications. Additionally, since Volume Shadow Copy snapshots are associated with the volume they belong to and Carbonite Migrate mirrors and replicates the data on the volume and not the volume itself, snapshots taken on the source cannot be used on the target's volume. Therefore, snapshots taken on the source are not mirrored or replicated to the target.

Note the following replication caveats.

1. FAT and FAT32 are not supported.
2. ReFS is only supported on Windows 2016 servers.
3. You must mirror and replicate to like file systems. For example, you cannot use NTFS to ReFS or ReFS to NTFS. You must use NTFS to NTFS or ReFS to ReFS. Additionally, you cannot have ReFS volumes mounted to mount points in NTFS volumes or NTFS volumes mounted to mount points in ReFS volumes.
4. You cannot replicate from or to a mapped drive.
5. If any directory or file contained in your job specifically denies permission to the system account or the account running the Double-Take service, the attributes of the file on the target will not be updated because of the lack of access. This also includes denying permission to the Everyone group because this group contains the system account.
6. If you select a dynamic volume and you increase the size of the volume, the target must be able to compensate for an increase in the size of the dynamic volume.
7. If you select files with alternate data streams, keep in mind the following.
 - a. Alternate data streams are not included in the job size calculation. Therefore, you may see the mirror process at 99-100% complete while mirroring continues.
 - b. The number of files and directories reported to be mirrored will be incorrect. It will be off by the number of alternate streams contained in the files and directories because the alternate streams are not counted. This is a reporting issue only. The streams will be mirrored correctly.
 - c. Use the file attributes and data comparison option when performing a difference mirror or verification to ensure that all alternate data streams are compared correctly.
 - d. If your alternate streams are read-only, the times may be flagged as different if you are creating a verification report only. Initiating a remirror with the verification will correct this issue.
8. If you select encrypted files, keep in mind the following.

- a. Only the data, not the attributes or security/ownership, is replicated. However, the encryption key is included. This means that only the person who created the encrypted file on the source will have access to it on the target.
 - b. Only data changes cause replication to occur; changing security/ownership or attributes does not.
 - c. Replication will not occur until the Windows Cache Manager has released the file. This may take awhile, but replication will occur when Carbonite Migrate can access the file.
 - d. When remirroring, the entire file is transmitted every time, regardless of the remirror settings.
 - e. Verification cannot check encrypted files because of the encryption. If remirror is selected, the entire encrypted file will be remirrored to the target. Independent of the remirror option, all encrypted files will be identified in the verification log.
 - f. Empty encrypted files will be mirrored to the target, but if you copy or create an empty encrypted file within the job after mirroring is complete, the empty file will not be created on the target. As data is added to the empty file on the source, it will then be replicated to the target.
 - g. When you are replicating encrypted files, a temporary file is created on both the source and target servers. The temporary file is automatically created in the same directory as the Carbonite Migrate disk queues. If there is not enough room to create the temporary file, an out of disk space message will be logged. This message may be misleading and indicate that the drive where the encrypted file is located is out of space, when it actually may be the location where the temporary file is trying to be created that is out of disk space.
 - h. Carbonite Migrate supports mirroring and replication of data stored on BitLocker enabled volumes when using the certificate based authentication method. Trusted Platform Module (TPM) is not supported because TPM uses a microchip that is built into the hardware to store the encryption keys. That same microchip would not be present on the target after failover.
9. If you are using mount points, keep in mind the following.
- a. By default, the mount point data will be stored in a directory on the target. You can create a mount point on the target to store the data or maintain the replicated data in a directory. If you use a directory, it must be able to handle the amount of data contained in the mount point.
 - b. Recursive mount points are not supported. If you select data stored on a recursive mount point, mirroring will never finish.
10. Carbonite Migrate supports transactional NTFS (TxF) write operations, with the exception of TxF SavePoints (intermediate rollback points).
- a. With transactional NTFS and Carbonite Migrate mirroring, data that is in a pending transaction is in what is called a transacted view. If the pending transaction is committed, it is written to disk. If the pending transaction is aborted (rolled back), it is not written to disk.
- During a Carbonite Migrate mirror, the transacted view of the data on the source is used. This means the data on the target will be the same as the transacted view of the data on the source. If there are pending transactions, the Carbonite Migrate **Target Data State** will indicate **Transactions Pending**. As the pending transactions are committed or aborted, Carbonite Migrate mirrors any necessary changes to the target. Once all pending transactions are completed, the **Target Data State** will update to **OK**.

If you see the pending transactions state, you can check the Carbonite Migrate log file for a list of files with pending transactions. As transactions are committed or aborted, the list is updated until all transactions are complete, and the **Target Data State** is **OK**.

- b. During replication, transactional operations will be processed on the target identically as they are on the source. If a transaction is committed on the source, it will be committed on the target. If a transaction is aborted on the source, it will be aborted on the target.
 - c. When cutover occurs any pending transactions on the target will be aborted.
11. Carbonite Migrate supports Windows symbolic links and junction points. A symbolic link is a link (pointer) to a directory or file. Junction points are links to directories and volumes.
- a. If the link and the file/directory/volume are both in your job, both the link and the file/directory/volume are mirrored and replicated to the target.
 - b. If the link is in the job, but the file/directory/volume it points to is not, only the link is mirrored and replicated to the target. The file/directory/volume that the link points to is not mirrored or replicated to the target. A message is logged to the Carbonite Migrate log identifying this situation.
 - c. If the file/directory/volume is in the job, but the link pointing to it is not, only the file/directory/volume is mirrored and replicated to the target. The link pointing to the file/directory/volume is not mirrored or replicated to the target.
 - d. Junction points that are orphans (no counterpart on the source) will be processed for orphan files, however, the contents of a junction point (where it redirects you) will not be processed for orphan files.
12. If you have the Windows NtfsDisable8dot3NameCreation setting enabled on the source but disabled on the target, there is a potential that you could overwrite and lose data on the target because of the difference in how long file names will be associated with short files names on the two servers. This is only an issue if there are like named files in the same directory (for example, longfilename.doc and longfi~1.doc in the same directory). To avoid the potential for any data loss, the NtfsDisable8dot3NameCreation setting should be the same on both the source and target.
13. Carbonite Migrate can replicate paths up to 32,760 characters, although each individual component (file or directory name) is limited to 259 characters. Paths longer than 32760 characters will be skipped and logged.
14. If you rename the root folder of a job, Carbonite Migrate interprets this operation as a move from inside the job to outside the job. Therefore, since all of the files under that directory have been moved outside the job and are no longer a part of the job, those files will be deleted from the target replica copy. This, in essence, will delete all of your replicated data on the target. If you have to rename the root directory of your job, make sure that the job is not connected.
15. Keep in mind the following caveats when including and excluding data for replication.
- a. Do not exclude Microsoft Word temporary files from your job. When a user opens a Microsoft Word file, a temporary copy of the file is opened. When the user closes the file, the temporary file is renamed to the original file and the original file is deleted. Carbonite Migrate needs to replicate both the rename and the delete. If you have excluded the temporary files from your job, the rename operation will not be replicated, but the delete operation will be replicated. Therefore, you will have missing files on your target.
 - b. When Microsoft SQL Server databases are being replicated, you should always include the tempdb files, unless you can determine that they are not being used by any application. Some applications, such as PeopleSoft and BizTalk, write data to the tempdb file. You can, most likely, exclude temporary databases for other database applications, but you should consult the product documentation or other support resources before doing so.


- c. Some applications create temporary files that are used to store information that may not be necessary to replicate. If user profiles and home directories are stored on a server and replicated, this could result in a significant amount of unnecessary data replication on large file servers. Additionally, the \Local Settings\Temporary Internet Files directory can easily reach a few thousand files and dozens of megabytes. When this is multiplied by a hundred users it can quickly add up to several gigabytes of data that do not need to be replicated.
 - d. Creating jobs that only contain one file may cause unexpected results. If you need to replicate just one file, add a second file to the job to ensure the data is replicated to the correct location. (The second file can be a zero byte file if desired.)
- 16. Carbonite Migrate does not replicate the last access time if it is the only thing that has changed. Therefore, if you are performing incremental or differential backups on your target machine, you need to make sure that your backup software is using an appropriate flag to identify what files have been updated since the last backup. You may want to use the last modified date on the file rather than the date of the last backup.
- 17. Keep in mind the following caveats when using anti-virus protection.
 - a. Virus protection software on the target should not scan replicated data. If the data is protected on the source, operations that clean, delete, or quarantine infected files will be replicated to the target by Carbonite Migrate. If the replicated data on the target must be scanned for viruses, configure the virus protection software on both the source and target to delete or quarantine infected files to a different directory that is not in the job. If the virus software denies access to the file because it is infected, Carbonite Migrate will continually attempt to commit operations to that file until it is successful, and will not commit any other data until it can write to that file.
 - b. You may want to set anti-virus exclusions on your source to improve replication performance. There are risks associated with making exclusions, so implement them carefully. For more information, see the Microsoft article [822158 Virus scanning recommendations for Enterprise computers that are running currently supported versions of Windows](#).
 - c. If you are using avast! anti-virus software, it must be installed in its default installation location if you want to protect your sever with a full server protection job. If it is not in its default installation directory, cutover will fail.
- 18. SQL Server 2005 or later may not initialize empty space when the database size increases due to the auto grow feature. Therefore, there is nothing for Carbonite Migrate to replicate when this empty space is created. When the empty space is populated with data, the data is replicated to the target. A verification report will report unsynchronized bytes between the source and target due to the empty space. Since the space is empty, the data on the source and target is identical. In the event of a failure, the SQL database will start without errors on the target.
- 19. If you have reparse points in your data set, Carbonite Migrate will replicate the tag, unless it is a known driver. If it is a known driver, for example Microsoft SIS, Carbonite Migrate will open the file allowing the reparse driver to execute the file. In this case, the entire file will be replicated to the target (meaning the file is no longer a reparse point on the target and has all the data).
- 20. Keep in mind if you have reparse points in your data set, the reparse driver cannot be loaded on the target during migration. You must load the reparse driver on the target after cutover in order to access the data. Additionally, you cannot have reparse points in your data set if you are using same server migration because the server is functioning as both a source and target.

21. If you are using the Microsoft Windows Update feature, keep in mind the following caveats.
 - a. Schedule your Windows Update outside the times when a mirroring operation (initial mirror or remirror) is running. In some cases, Windows Update may perform an NTFS transactional rollback before displaying the dialog box to reboot the computer. This rollback will cause a mirror. If that mirror completes before the reboot, the reboot will trigger another mirror, unless you have configured Carbonite Migrate to only mirror changed files on reboot.
 - b. You must resolve any Windows Update incomplete operations or errors before cutover. (Check the windowsupdate.log file.) Also, do not cutover if the target is waiting on a Windows Update reboot. If cutover occurs before the required Windows Update reboot, the target may not operate properly or it may not boot. You could also get into a situation where the reboot repeats indefinitely. One possible workaround for the reboot loop condition is to access a command prompt through the Windows Recovery Environment and delete the file \Windows\winsxs\pending.xml file. You may need to take ownership of the file to delete it. Contact technical support for assistance with this process or to evaluate other alternatives. Before you contact technical support, you should use the Microsoft System Update Readiness Tool as discussed in [Microsoft article 947821](#). This tool verifies and addresses many Windows Update problems.
22. If you are using Windows deduplication, keep in mind the following caveats.
 - a. Deduplicated data on the source will be expanded to its original size on the target when mirrored. Therefore, you must have enough space on the target for this expansion, even if you have deduplication enabled on the target.
 - b. If you have deduplicated data on the target, mirroring and replication (like any other write process) will create a new file or new blocks of data. Existing blocks of deduplicated data will remain as they were until the next garbage collection.
 - c. If you are protecting an entire server, you must have the deduplication feature installed on both the source and target. It can be enabled or disabled independently on the two servers, but it must at least be installed on both of the servers.
 - d. After cutover, the amount of disk space on the failed over server will be incorrect until you run the deduplication garbage collection which will synchronize the disk space statistics.
23. Replication is not case-sensitive. For example, if you rename the file Test.txt to test.txt, that change will not be replicated to the target. You will have to delete the file on the target and when it is mirrored, the new case of the file name will be used.
24. If you are using Windows 2008 R2, virtual hard disks can be mounted and dismounted reusing the same drive letter. However, once you have established a job, you cannot mount a different virtual hard disk to the same drive letter used in your job. This could cause errors, orphan files, or possibly data corruption. If you must change drive letters associated with a virtual hard disk, delete the job, change the mounting, and then re-create the job.

Chapter 3 Carbonite Replication Console

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

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



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



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

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

Carbonite Replication Console requirements

You must meet the following requirements for the Carbonite Replication Console.

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



The Carbonite Migrate 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 Carbonite Migrate while running in full operating system mode. In any case, you cannot run the Carbonite Replication Console on Server Core.

Console options

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

- **Monitoring**—This section is used to determine how the console monitors your Carbonite Migrate 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 Carbonite Replication 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 Carbonite Migrate servers.
 - **Default port for XML web services protocol**—Specifies the port that the console will use when sending and receiving data to Carbonite Migrate 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 Carbonite Migrate version, you will need to use the legacy protocol port. This applies to Carbonite Migrate 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 Carbonite Replication Console. Use this button as directed by technical support.
 - **View Log File**—This button opens the Carbonite Replication Console log file. Use this button as directed by technical support. You can also select **View, View Console Log File** to open the Carbonite Replication Console log file.
 - **View Data File**—This button opens the Carbonite Replication Console data file. Use this button as directed by technical support. You can also select **View, View Console Data File** to open the Carbonite Replication Console data file.
- **License Inventory**—This section controls if the console contains a license inventory. This

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

- **Enable license inventory**—This option allows you to use this console to manage the Carbonite Migrate licenses assigned to your organization. When this option is enabled, the **License Inventory** page is also enabled.
- **Default Installation Options**—All of the fields under the **Default Installation Options** section are used by the push installation on the **Install** page. The values specified here will be the default options used for the push installation.
 - **Activate online after install completes**—Specify if you want to activate your Carbonite Migrate licenses at the end of the installation. The activation requires Internet access from the console machine or the machine you are installing to. Activation will be attempted from the console machine first and if that fails, it will be attempted from the machine you are installing to. If you choose not to have the installation activate your licenses, you will have to activate them through the console license inventory or the server's properties page.
 - **Location of install folders**—Specify the parent directory location where the installation files are located. The parent directory can be local on your console machine or a UNC path.
 - **Windows**—Specify the parent directory where the Windows installation file is located. The default location is where the Carbonite Replication Console is installed, which is `\Program Files\Carbonite\Replication`. The console will automatically use the `\x64` subdirectory which is populated with the Windows installation files when you installed the console. If you want to use a different location, you must copy the `\x64` folder and its installation file to the different parent directory that you specify.
 - **Linux**—Specify the parent directory where the Linux installation files are located. The default location is where the Carbonite Replication Console is installed, which is `\Program Files\Carbonite\Replication`. The console will automatically use the `\Linux` subdirectory, however that location will not be populated with the Linux installation files when you installed the console. You must copy the Linux `.deb` or `.rpm` files from your download to the `\Linux` subdirectory in your Carbonite Replication Console installation location. Make sure you only have a single version of Linux installation files. The push installation cannot determine which version to install if there are multiple versions in the `\Linux` subdirectory. If you want to use a different location, you must copy the `\Linux` folder and its installation files to the different parent directory that you specify.
- **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 Carbonite Migrate where the installation files will be copied and run.
 - **Installation folder**—Specify the location where you want to install Carbonite Migrate on each server. This field is not used if you are upgrading an existing version of Carbonite Migrate. In that case, the existing installation folder will be used.
 - **Queue folder**—Specify the location where you want to store the Carbonite Migrate disk queue on each server.
 - **Amount of system memory to use**—Specify the maximum amount of memory, in MB, that can be used for Carbonite Migrate 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, Carbonite Migrate will automatically begin the auto-disconnect process.
 - **Unlimited disk queue**—Carbonite Migrate 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, Carbonite Migrate 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 Carbonite Migrate disk queuing. When the disk space limit is reached, Carbonite Migrate 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 Carbonite Replication 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 Carbonite Migrate where the installation files will be copied and run.

Chapter 4 Managing servers

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

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

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



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

Left pane

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

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

Create New Server Group

Creates a new server group below the selected group

Rename Server Group

Allows you to rename the selected server group

Delete Server Group

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

Overflow Chevron








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

Top right pane

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





Column 1 (Blank)

The first blank column indicates the machine type.

-  Carbonite Migrate source or target server which could be a physical server, virtual machine, or a cluster node
-  Carbonite Migrate source or target server which is a Windows cluster
-  vCenter server
-  ESX server
-  Carbonite Migrate Reporting Service server
-  Offline server which means the console cannot communicate with this machine.
-  Any server icon with a red circle with white X overlay is an error which means the console can communicate with the machine, but it cannot communicate with Carbonite Migrate 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 32.

Version

The Carbonite Migrate product version information, if any.

Licensing Status

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

Product

The Carbonite Migrate products, if any, licensed for the server

Bottom right pane

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

Name

The name or IP address of the server.

Operating system

The operating system of the server. This field will not be displayed if the console cannot connect to Carbonite Migrate on the server.

Product

The Carbonite Migrate products, if any, licensed for the server

Version

The product version information, if any

Serial Number

The serial number associated with the Carbonite Migrate license

Toolbar

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

Create a New Job

The available job creation choices depend on the Carbonite Migrate licenses applied to your server.

- **Protect**—If you are licensed for Carbonite Availability, use the **Protect** option to create a protection job for the selected server.
- **Migrate**—If you are licensed for Carbonite Migrate or certain Carbonite Availability licenses, use the **Migrate** option to create a migration job for the selected server.

Add Servers

Adds a new server. This button leaves the **Servers** page and opens the **Add Servers** page. See *Adding servers* on page 28.

View Server Details

Views detailed information about a server. This button leaves the **Servers** page and opens the **View Server Details** page. See *Viewing server details* on page 32.

Edit Server Properties

Edits the server's properties and options. This button leaves the **Servers** page and opens the **Edit Server Properties** page. See *Editing server properties* on page 34.

Remove Server

Removes the server from the console.

Provide Credentials

Changes the login credentials that the Carbonite Replication 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 31. You will remain on the **Servers** page after updating the server credentials. If your jobs use the same credentials, make sure you also update the credentials for any active jobs on the server. See the *Managing and controlling* section for your specific job type.

If you are using domain credentials for your Carbonite Migrate servers and you change those credentials, you will continue to receive a Windows Security pop-up in the Carbonite Replication 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 Carbonite Migrate servers using **Provide Credentials** in order to terminate the repeated pop-ups.

Manage Group Assignments

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

Install

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

Uninstall

Uninstalls Carbonite Migrate on the selected server.

View Server Events

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

View Server Logs

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

Launch Reporting

Launches the Reporting Service report viewer.

Activate Online

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

Refresh

Refreshes the status of the selected servers.

Search

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

Overflow Chevron

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

Right-click menu

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

Protect

If you are licensed for Carbonite Availability, use the **Protect** option to create a protection job for the selected server.

Migrate

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

View Server Details

Views detailed information about a server. This button leaves the **Servers** page and opens the **View Server Details** page. See *Viewing server details* on page 32.

Edit Server Properties

Edits the server's properties and options. This button leaves the **Servers** page and opens the **Edit Server Properties** page. See *Editing server properties* on page 34.

Remove Server

Removes the server from the console.

Provide Credentials

Changes the login credentials that the Carbonite Replication 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 31. You will remain on the **Servers** page after updating the server credentials. If your jobs use the same credentials, make sure you also update the credentials for any active jobs on the server. See the *Managing and controlling* section for your specific job type.

If you are using domain credentials for your Carbonite Migrate servers and you change those credentials, you will continue to receive a Windows Security pop-up in the Carbonite Replication 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 Carbonite Migrate servers using **Provide Credentials** in order to terminate the repeated pop-ups.

Manage Group Assignments

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

Install

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

Uninstall

Uninstalls Carbonite Migrate on the selected server.

Copy

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

Paste

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

View Server Events

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

View Server Logs

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

Launch Reporting

Launches the Reporting Service report viewer.

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 Carbonite Migrate log files; Carbonite Migrate 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 Carbonite Replication Console while monitoring the replication service details. You can open multiple **Replication service view** windows for multiple servers. When the Carbonite Replication 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 Carbonite Replication Console from within the **Replication service view** window. See the *Reference Guide* for a complete list of replication details.

Refresh

Refreshes the status of the selected servers.

Adding servers

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

Inserting servers manually

1. Click **Servers** from the main toolbar.
2. Click **Add servers** from the **Servers** page toolbar.
3. On the **Manual Entry** tab, specify the server information.
 - **Server**—This is the name or IP address of the server or appliance to be added to the console.



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

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

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



If you are using domain credentials for your Carbonite Migrate servers and you change those credentials, you will continue to receive a Windows Security pop-up in the Carbonite Replication 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 Carbonite Migrate servers in the Carbonite Replication Console **Servers** page in order to terminate the repeated pop-ups.

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

6. If you need to remove servers or appliances from the list of **Servers to be added**, highlight a server and click **Remove**. You can also remove all of them with the **Remove All** button.
7. When your list of **Servers to be added** is complete, click **OK**.

Inserting servers through Active Directory discovery

You can insert servers using Active Directory discovery.

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

Importing and exporting servers from a server and group configuration file

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

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

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

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

Providing server credentials

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



If you are using domain credentials for your Carbonite Migrate servers and you change those credentials, you will continue to receive a Windows Security pop-up in the Carbonite Replication 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 Carbonite Migrate servers in the Carbonite Replication Console in order to terminate the repeated pop-ups.

Viewing server details

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

Server name

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

Operating system

The server's operating system version

Roles

The role of this server in your Carbonite Migrate environment. In some cases, a server can have more than one role.

- **Engine Role**—Source or target server
- **Reporting Service**—Reporting Service server

Status

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

Activity

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

Connected via

The IP address and port the server is using for communications. You will also see the Carbonite Migrate protocol being used to communicate with server. The protocol will be XML web services protocol (for servers running Carbonite Migrate 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

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

- *General server properties* on page 35—Identifies the server
- *Server licensing* on page 36—Views, adds, and removes license keys
- *Server setup properties* on page 38—Indicates how the server will act on startup and shutdown
- *Carbonite Migrate queue* on page 41—Configures the Carbonite Migrate queues
- *Source server properties* on page 45—Configures the source server
- *Target server properties* on page 47—Configures the target server
- *E-mail notification configuration* on page 49—Configures e-mail notification
- *Script credentials* on page 51—Specifies credentials to be used when executing custom scripts during mirroring or cutover
- *Log file properties* on page 52—Configures log files

General server properties

The general server properties identify the server.



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

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

- **Default address**—On a server with multiple NICs, you can specify which address Carbonite Migrate 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 Carbonite Migrate servers. If you change the port, you must stop and restart the Double-Take service.
- **Encrypt network data**—Use this option to encrypt your data before it is sent from the source to the target. Both the source and target must be encryption capable (version 7.0.1 or later), however this option only needs to be enabled on the source or target in order to encrypt data. Keep in mind that all jobs from a source with this option enabled or to a target with this option enabled will have the same encryption setting. Changing this option will cause jobs to auto-reconnect and possibly remirror. The encryption method used is AES-256.
- **Automatic retry**—This option will have the target server automatically retry server login credentials for a job, after the specified retry interval, if the server login credentials are not accepted. Keep in mind the following caveats when using this option.
 - Because server logins for a job are controlled by the target, this setting is only applicable to target servers.
 - This is only for server credentials, not job credentials.
 - Verify your environment's security policy when using this option. Check your policies for failed login lock outs and resets. For example, if your policy is to reset the failed login attempt count after 30 minutes, set this auto-retry option to the same or a slightly larger value as the 30 minute security policy to decrease the chance of a lockout.
- **Retry on this interval**—If you have enabled the automatic retry, specify the length of time, in minutes, to retry the login.

Server licensing

Licensing identifies your Carbonite Migrate license keys.



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

	Product	Serial Number	Expiration Date	Licen
⚠	DoubleTake Availability Virtual	4567	12/26/2017	knc0-

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



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

- **Current license keys**—The server's current license key information is displayed. To remove a key, highlight it and click **Remove**. To copy a key, highlight it and click **Copy**. To replace a key, enter a new key and click **Add**. If you are replacing an unexpired key with the same version and serial number, you should not have to reactivate it and any existing jobs will continue uninterrupted. If you are replacing an unexpired key with a new version or new serial number or

replacing an expired key, you will have to reactivate and remirror.

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



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

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


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



For Carbonite Migrate, license keys do not have a grace period and must be activated in order to be used. Once the license has been activated, you will have a specific number of days (generally 30 days) to complete your migration process, depending on your license type.

- **Metered Usage Licensing**—If you are a service provider participating in the Metered Usage program, you can configure the metered usage license for your target servers here. If you are not in this program, you can skip this section. For the latest and complete details on Metered Usage, see the help link in the metered usage web portal.
 1. Specify your **Service provider account number**. The account number is displayed in the upper right corner of the metered usage portal.
 2. Specify the **Customer name**. Use the customer name configured on the Customers list in the metered usage portal.
 3. Select the appropriate **Product** that corresponds with the Carbonite Migrate product being used.
 4. Click **Submit** to activate the metered usage license on the target.

Server setup properties



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

- **Log statistics automatically**—If enabled, Carbonite Migrate statistics logging will start automatically when Carbonite Migrate is started.
- **Enable task command processing**—Task command processing is a Carbonite Migrate 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 Carbonite Migrate 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**—For servers in a domain, if this option is enabled, the Double-Take service registers with Windows Active Directory when the service is started.
- **Mirror only changed files when source reboots**—If enabled, Carbonite Migrate will use the Carbonite Migrate driver change journal to track file changes. If the source is rebooted, only the files identified in the change journal will be remirrored to the target. This setting helps improve mirror times. If this option is enabled but the change journal 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 Carbonite Migrate 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, Carbonite Migrate will automatically disconnect jobs when necessary. If this option is enabled, Carbonite Migrate 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, Carbonite Migrate will automatically reconnect any jobs that were previously connected. Then, if configured, Carbonite Migrate 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, Carbonite Migrate will automatically start disconnecting jobs. This is called auto-disconnect. The transaction logs and system memory are flushed allowing Carbonite Migrate to begin processing anew. The auto-reconnect process ensures that any jobs that were auto-disconnected are automatically reconnected. Then, if configured, Carbonite Migrate 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 Carbonite Migrate will auto-reconnect. If configured, Carbonite Migrate 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, Carbonite Migrate 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, Carbonite Migrate 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 Carbonite Migrate source, not being able to communicate with the Carbonite Migrate 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 Carbonite Migrate is able to successfully persist this information, when the Double-Take service on the target is restarted, Carbonite Migrate will pick up where it left off, without requiring an auto-disconnect, auto-reconnect, or auto-remirror. If Carbonite Migrate 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,

Carbonite Migrate 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 Carbonite Migrate 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, Carbonite Migrate 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 Carbonite Migrate will perform the mirror when it is automatically remirroring.

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.**—Carbonite Migrate 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 option requires no time for comparison, but it can be slower because it sends the entire file. However, it is useful for configurations that have large data sets with millions of small files that are frequently changing and it is more efficient to send the entire file. You may also need to use this option if configuration management policies require sending the entire file.
- **Compare file attributes. Send the entire file.**—Carbonite Migrate will compare file attributes and will mirror those files that have different attributes, sending the entire file. This option is the fastest comparison method, but can be slower because it sends the entire file. However, it is useful for configurations that have large data sets with millions of small files that are mostly static and not changing. You may also need to use this option if configuration management policies require sending the entire file. This option is not available for full server jobs.
- **Compare file attributes. Send the attributes and bytes that differ.**—Carbonite Migrate will compare file attributes and will mirror only the attributes and bytes that are different. This option is the fastest comparison method and fastest mirror option. Files that have not changed can be easily skipped. Also files that are open and require a checksum mirror can be compared.
- **Compare file attributes and data. Send the attributes and bytes that differ.**—Carbonite Migrate will compare file attributes and the file data and will mirror only the attributes and bytes that are different. This comparison method is not as fast because every file is compared, regardless of whether the file has changed or is open. However, sending only the attributes and bytes that differ is the fastest mirror option.
- **Time allowed to complete shutdown operations**—This setting indicates the amount of time, in seconds, for the Double-Take service to wait prior to completing a shutdown so that Carbonite Migrate 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 Carbonite Migrate. It does not control the service shutdown through a reboot or from the Service Control Manager.

Carbonite Migrate queue

During the Carbonite Migrate installation, you identified the amount of disk space that can be used for Carbonite Migrate queuing. Queuing to disk allows Carbonite Migrate 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.

Carbonite Migrate 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 Carbonite Migrate 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 Carbonite Migrate 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, Carbonite Migrate 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. Carbonite Migrate 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 Carbonite Migrate 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, Carbonite Migrate pulls the data from system memory. Carbonite Migrate 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 Carbonite Migrate server.

The screenshot shows the 'Queue' configuration dialog box. It has a title bar with an up arrow and the word 'Queue'. Below the title bar, there is a 'Queue folder:' label followed by a text box containing 'C:\Program Files\Carbonite\Replication\'. To the right of the text box is a browse button with three dots. Below the text box is a note: 'A dedicated non-boot volume will offer the best performance and reliability.' There are two columns of settings. The first column has 'Amount of system memory to use (MB):' with a spin box set to '1024'. Below it are three radio buttons: 'Do not use disk queue', 'Unlimited disk queue' (which is selected), and 'Limit disk space for queue (MB):' with a spin box set to '0'. The second column has 'Minimum free disk space (MB):' with a spin box set to '250'. Below it is 'Alert at this queue usage (%):' with a spin box set to '50'. At the bottom of the dialog, there is a note: 'Changing the queue location or memory usage will not take effect until the service has been restarted.'

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

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

- Select an NTFS volume. Do not select a ReFS volume.
- 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 Carbonite Migrate 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 Carbonite Migrate queues.

- **Amount of system memory to use**—This is the maximum amount of Windows system memory, in MB, that Carbonite Migrate 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, Carbonite Migrate will use less system memory, but you will queue to disk sooner which may impact system performance. If you set it higher, Carbonite Migrate 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 Carbonite Migrate 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 Carbonite Migrate 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 Carbonite Migrate, you might configure the application to use 1 GB and Carbonite Migrate 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, Carbonite Migrate will automatically begin the auto-disconnect process.
- **Unlimited disk queue**—Carbonite Migrate 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, Carbonite Migrate 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 Carbonite Migrate disk queuing. When the disk space limit is reached, Carbonite Migrate 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, Carbonite Migrate will only be able to use 4.5 GB so that 500 MB remains free.

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

Source server properties

These properties are specific to the source server role.

Source

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

Replicate NTFS security attributes by name

Maximum pending mirror operations:
1000

Size of mirror packets (bytes):
65536

Calculate size of protected data upon connection

- **Number of replication packets per one mirror packet**—You can specify the ratio of replication packets to mirror packets that are placed in the source queue. The default value (5) allows Carbonite Migrate 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, Carbonite Migrate 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 Carbonite Migrate 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 Carbonite Migrate 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 Carbonite Migrate 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 Carbonite Migrate to queue more data for transfer.
- **Size of mirror packets**—This option determines the size of the mirror packets, in bytes, that Carbonite Migrate 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 Carbonite Migrate 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.

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 Carbonite Migrate processing. After cutover, any target paths that are blocked will be unblocked automatically during the cutover process so that users can modify data after cutover.



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.

- **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. Carbonite Migrate 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, Carbonite Migrate 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.

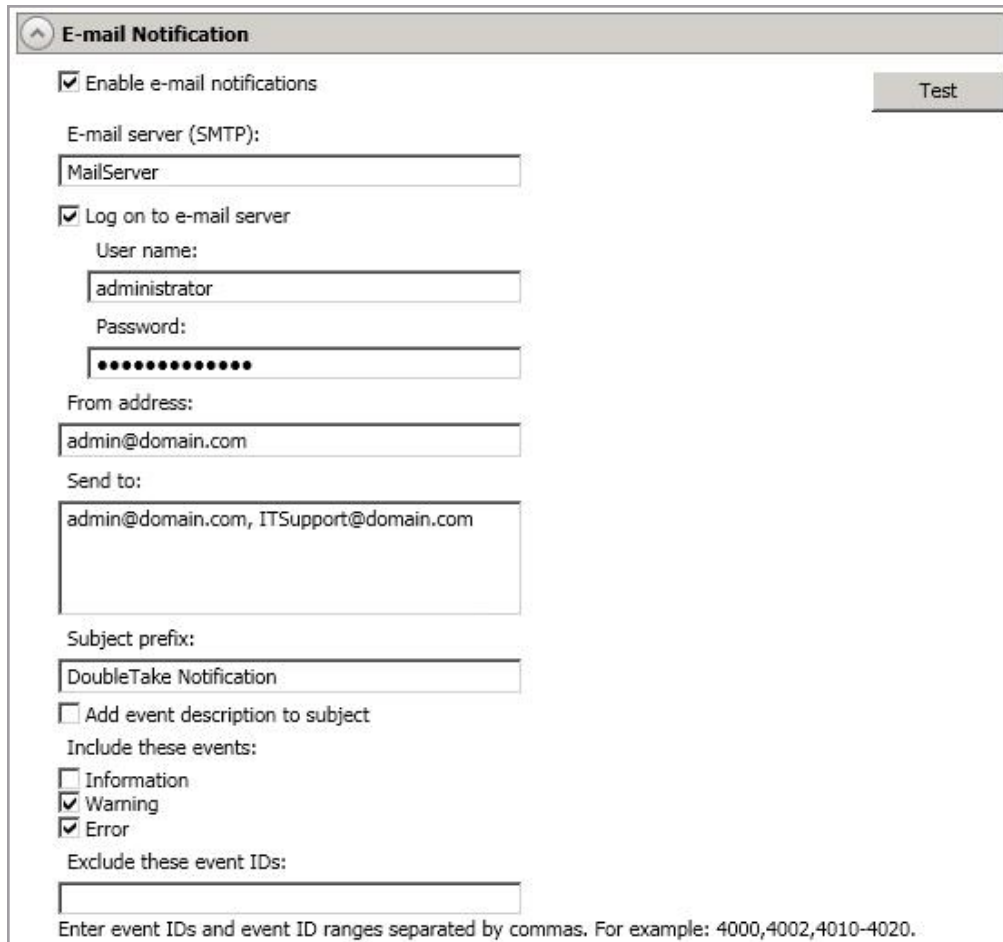
Alternate data streams that are deleted on the source will not be moved on the target.

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 Carbonite Migrate event messages to specific addresses using an SMTP mail server. The subject of the e-mail will contain an optional prefix, the server name where the message was logged, the message ID, and the severity level (information, warning, or error). The text of the event message will be displayed in the body of the e-mail message.



The screenshot shows a configuration window titled "E-mail Notification". It contains the following fields and options:

- Enable e-mail notifications (with a "Test" button to the right)
- E-mail server (SMTP): MailServer
- Log on to e-mail server
 - User name: administrator
 - Password: [masked]
- From address: admin@domain.com
- Send to: admin@domain.com, ITSupport@domain.com
- Subject prefix: DoubleTake Notification
- Add event description to subject
- Include these events:
 - Information
 - Warning
 - Error
- Exclude these event IDs: [empty field]

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

- **Enable e-mail notification**—This option enables the e-mail notification feature. Any specified notification settings will be retained if this option is disabled.
- **E-mail server**—Specify the name of your SMTP mail server.
- **Log on to e-mail server**—If your SMTP server requires authentication, enable this option and specify the **User name** and **Password** to be used for authentication. Your SMTP server must support the LOGIN authentication method to use this feature. If your server supports a different authentication method or does not support authentication, you may need to add the Carbonite Migrate 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 Carbonite Migrate e-mail message. The address is limited to 256 characters.

- **Send to**—Specify the e-mail addresses that each Carbonite Migrate e-mail message should be sent to. Enter the addresses as a comma or semicolon separated list. Each address is limited to 256 characters. You can add up to 256 e-mail addresses.
- **Subject prefix** and **Add event description to subject**—The subject of each e-mail notification will be in the format Subject Prefix : Server Name : Message Severity : Message ID : Message Description. The first and last components (Subject Prefix and Message Description) are optional. The subject line is limited to 100 characters.

If desired, enter unique text for the **Subject prefix** which will be inserted at the front of the subject line for each Carbonite Migrate e-mail message. This will help distinguish Carbonite Migrate 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 Carbonite Migrate e-mail messages are not blocked.

Script credentials

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



Script Credentials

These credentials will be used when executing custom scripts, such as during mirroring or failover.

User name:
administrator

Password:
●●●●●●●●●●

Domain:
domain.com

Test

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

Log file properties

These settings allow you to specify your log file configuration.

The screenshot shows a 'Logging' configuration window with the following sections and values:

- Logging folder:** C:\Program Files\Carbonite\Replication\
- Messages & Alerts:**
 - Maximum size (bytes): 5242880
 - Maximum number of files: 20
- Verification:**
 - File name: DTVerify.log
 - Maximum size (bytes): 1048576
 - Append
 - Language: English (United States)
- Statistics:**
 - File name: statistic.sts
 - Maximum size (bytes): 10485760
 - Write interval (minutes): 5

- **Logging folder**—Specify the directory where each of the log files in this section are stored. The default location is the directory where the Carbonite Migrate 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 54.
 - **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, Carbonite Migrate 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: 1/24/2018 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 = 1/17/2018 8:21:36 PM Size = 1272 Mask = [0x20]
  Target Attributes: Timestamp = 1/17/2018 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 = 1/17/2018 8:21:37 PM Size = 1272 Mask = [0x20]
  Target Attributes: Timestamp = 1/17/2018 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 = 1/12/2018 3:28:20 PM Size = 17 Mask = [0x20]
  Target Attributes: Timestamp = 1/20/2018 5:05:26 PM Size = 2 Mask = [0x20]
  17 BYTES OUT OF SYNC
Completion Time: 1/24/2018 12:37:44 PM for connection 2 (Sales data for alpha -->
206.31.65.40 : 1100)
Elapsed Time (seconds): 1320.256470
Total Directories Compared: 657
Total Directories Missing: 0
Total Directories Remirrored: 0
Total Files Compared: 120978
Total Files Missing: 0
Total Files Different: 3
Total Files Encrypted: 0
Total Files Remirrored: 1
Total Bytes Skipped: 0
Total Bytes Compared: 18527203678
Total Bytes Missing: 0
Total Bytes Different: 17
Total Bytes Remirrored: 17
Related links and directory attributes have been adjusted.
----- END OF VERIFICATION -----
```

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

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

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



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

Viewing server events

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

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


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

Viewing server logs

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

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

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



Time	Description	Service
6/22/2017 2:42:55 PM	C:\ --> Exc, Rec (2)	Double-Take
6/22/2017 2:42:55 PM	C:\PerfLogs --> Inc, Rec (2)	Double-Take
6/22/2017 2:42:55 PM	Service has established a connection to (172.31.206.201:6320) for Replication Set FilesAn...	Double-Take
6/22/2017 2:42:55 PM	Compression has been disabled (2)	Double-Take
6/22/2017 2:42:55 PM	ConnectionCmd succeeded: conid = 2	Management Service
6/22/2017 2:42:55 PM	Adding Connection: Id = 92a61fa2-387a-4759-9fc8-60bc55141f08, ReplicationSetName = ...	Management Service
6/22/2017 2:42:55 PM	Issuing MirrorCmd for conid 2 because a difference mirror was requested.	Management Service
6/22/2017 2:42:55 PM	MirrorCmd succeeded for conid 2	Management Service
6/22/2017 2:42:55 PM	Mirror Started, Differences, Fast Attribute, Block Checksum, Checksum All <2>	Double-Take
6/22/2017 2:42:55 PM	Calculating repset size... ConID: 2	Double-Take
6/22/2017 2:42:55 PM	The following (1) non-overlapping paths will be considered for mirroring	Double-Take
6/22/2017 2:42:55 PM	C:\PerfLogs	Double-Take
6/22/2017 2:42:55 PM	Only using the change journal re-mirror after a reboot.	Double-Take
6/22/2017 2:42:55 PM	---RepSet size calc done !!! Size: 0 Time spent: 94ms RepSet: FilesAndFolders_6d2bfd96a...	Double-Take
6/22/2017 2:42:57 PM	Processing path C:\PerfLogs	Double-Take
6/22/2017 2:42:57 PM	Mirror Ended <2>	Double-Take
6/22/2017 2:42:57 PM	Directories Mirrored: 0	Double-Take
6/22/2017 2:42:57 PM	Files Mirrored: 0	Double-Take
6/22/2017 2:42:57 PM	Bytes Mirrored: 0	Double-Take
6/22/2017 2:42:57 PM	Files Checksummed: 0	Double-Take
6/22/2017 2:42:57 PM	Elapsed Time: 2.511 seconds	Double-Take
6/22/2017 2:42:57 PM	Paused Time: 0.000 seconds	Double-Take
6/22/2017 2:42:57 PM	Total number of pauses: 0	Double-Take
6/22/2017 2:43:01 PM	Share query polling interval is invalid or not specified -- using default share query polling in...	Management Service
6/22/2017 2:43:03 PM	Starting monitor e8f5d8cb-e09a-43b9-ab29-162548d2262d: name = ShareMonitor-6d2bfd96a...	Management Service
6/22/2017 2:43:03 PM	Failed activity on monitor named ShareMonitor-6d2bfd96a93-4eb4-8530-ca317cc7fcf9 (id...	Management Service
6/22/2017 2:43:03 PM	No action specified.	Management Service
6/22/2017 2:43:03 PM	Triggered actions complete for monitor named ShareMonitor-6d2bfd96a93-4eb4-8530-ca...	Management Service
6/22/2017 2:43:03 PM	Persisting shares for job 6d2bfd96a93-4eb4-8530-ca317cc7fcf9	Management Service
6/22/2017 2:43:33 PM	The monitor named ShareMonitor-6d2bfd96a93-4eb4-8530-ca317cc7fcf9 (id = e8f5d8cb-...	Management Service
6/22/2017 2:59:24 PM	Service has stopped Replication to BETA (172.31.206.201:6320) for Replication Set FilesAn...	Double-Take
6/22/2017 2:59:24 PM	Removing all repsets for source BETA ip://172.31.206.201:6320 in Target Status Map.	Double-Take
6/22/2017 2:59:24 PM	Updating connection status for source ip://172.31.206.201:6320, target ip://172.31.206.2...	Double-Take
6/22/2017 2:59:24 PM	Source lost connection with BETA ip://172.31.206.201:6320	Double-Take
6/22/2017 2:59:24 PM	Removing BETA ip://172.31.206.201:6320 from auto-reconnect list.	Double-Take
6/22/2017 2:59:24 PM	Service has disconnected from BETA (172.31.206.201:6320) for Replication Set FilesAndFol...	Double-Take
6/22/2017 2:59:24 PM	Disable replication on drive: \Device\HarddiskVolume2 with path: C:\	Double-Take
6/22/2017 2:59:24 PM	StopReplicationHandler set state to IDLE	Double-Take
6/22/2017 2:59:35 PM	Deleting monitor e8f5d8cb-e09a-43b9-ab29-162548d2262d: name = ShareMonitor-6d2bfd96a...	Management Service

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

Start 

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

Pause 

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

Copy 

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

Clear 

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

Filter

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

Managing VMware servers

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

VMware Server

The name of the VMware server

Full Name

The full name of the VMware server

User Name

The user account being used to access the VMware server

Add VMware Server

Add a new VMware server. When prompted, specify the VMware server and a user account. If you are using a non-default port for your server, specify the server followed by a colon and then the port number, for example, 112.47.12.7:85. If your server name does not match the security certificate or the security certificate has expired, you will be prompted if you want to install the untrusted security certificate.

Remove Server

Remove the VMware server from the console.

Provide Credentials

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

Chapter 5 Files and folders migration

Create a files and folders migration job when you want to migrate data or file shares. This job type does not migrate a server's system state.

- *Files and folders migration requirements* on page 61—Files and folders migration includes specific requirements for this type of migration.
- *Creating a files and folders migration job* on page 65—This section includes step-by-step instructions for creating a files and folders migration job.
- *Managing and controlling files and folders migration jobs* on page 83—You can view status information about your files and folders migration job.
- *Cutting over files and folders migration jobs* on page 102—Use this section when you are ready to cutover from your source to your target, which contains the data you migrated from your source.

Files and folders migration requirements

Use the following requirements for files and folders migration.

- **Operating system**—The following operating systems are supported for files and folders migrations jobs.
 - Windows 2016 and Server Core 2016
 - Windows 2012 R2 and Server Core 2012 R2
 - Windows 2012 and Server Core 2012
 - Windows 2008 R2 Service Pack 1 or later and Server Core 2008 R2 Service Pack 1 or later



Windows 2016 support is for the primary operating system features available in Windows 2012. New operating system features specific to Windows 2016, such as Nano Server, Windows Containers, and so on, are not supported.

Due to VMware limitations, Windows 2016 is not supported on ESX 5.5.

DNS updates are not supported for Server Core servers.

-
- **Cloud**—Carbonite Migrate can be used to migrate data to an existing server in the cloud. Keep in mind that you should enable appropriate security measures, like VPN, to protect your data as it migrates to the cloud.
 - **File system**—Carbonite Migrate supports the NTFS file system. On Windows 2016 servers, ReFS is also supported. FAT and FAT32 are not supported. For detailed information on other file system capabilities, see *Mirroring and replication capabilities* on page 7.
 - **Microsoft .NET Framework**—Microsoft .NET Framework version 4.5.1 or later is required on the source and target.
 - **System memory**—The minimum system memory on each server is 1 GB.
 - **Disk space for program files**—This is the amount of disk space needed for the Carbonite Migrate program files. The amount depends on your operating system version and ranges from 350-500 MB.



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

Make sure you have additional disk space for Carbonite Migrate queuing, logging, and so on.

-
- **Server name**—Carbonite Migrate includes Unicode file system support, but your server name must still be in ASCII format. Additionally, all Carbonite Migrate servers and appliances must have a unique server name.
 - **Time**—The clock on your Carbonite Migrate servers must be within a few minutes of each other, relative to UTC. Large time skews (more than five minutes) will cause Carbonite Migrate errors.
 - **Protocols and networking**—Your servers must meet the following protocol and networking

requirements.

- Your servers must have TCP/IP with static IP addressing.
- IPv4 only configurations are supported, IPv4 and IPv6 are supported in combination, however IPv6 only configurations are not supported
- If you are using IPv6 on your servers, your clients must be run from an IPv6 capable machine.
- In order to properly resolve IPv6 addresses to a hostname, a reverse lookup entry should be made in DNS.
- If you are using Carbonite Migrate over a WAN and do not have DNS name resolution, you will need to add the host names to the local hosts file on each server running Carbonite Migrate.
- **NAT support**—Carbonite Migrate supports IP and port forwarding in NAT environments with the following caveats.
 - Only IPv4 is supported.
 - Only standalone servers are supported.
 - Make sure you have added your server to the Carbonite Replication Console using the correct public or private IP address. The name or IP address you use to add a server to the console is dependent on where you are running the console. Specify the private IP address of any servers on the same side of the router as the console. Specify the public IP address of any servers on the other side of the router as the console.
 - DNS failover and updates will depend on your configuration
 - Only the source or target can be behind a router, not both.
 - The DNS server must be routable from the target
- **Reverse lookup zone**—If you are using a DNS reverse lookup zone, then it must be Active Directory integrated. Carbonite Migrate is unable to determine if this integration exists and therefore cannot warn you during job creation if it doesn't exist.
- **DNS updates**—You can 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. Carbonite Migrate 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.

- **Windows firewall**—If you have Windows firewall enabled on your servers, there are two

requirements for the Windows firewall configuration.

- The Carbonite Migrate installation program will automatically attempt to configure ports 6320, 6325, and 6326 for Carbonite Migrate. If you cancel this step, you will have to configure the ports manually.
- If you are using the Carbonite Replication Console to push installations out to your Windows servers, you will have to open firewall ports for WMI (Windows Management Instrumentation), which uses RPC (Remote Procedure Call). By default, RPC will use ports at random above 1024, and these ports must be open on your firewall. RPC ports can be configured to a specific range by specific registry changes and a reboot. See the [Microsoft Knowledge Base article 154596](#) for instructions. Additionally, you will need to open firewall ports for SMB (server message block) communications which uses ports 135-139 and port 445, and you will need to open File and Printer Sharing. As an alternative, you can disable the Windows firewall temporarily until the push installations are complete.

See *Firewalls* on page 255 for instructions on handling firewalls in your environment.

- **Windows Management Instrumentation (WMI)**—Carbonite Migrate is dependent on the WMI service. If you do not use this service in your environment, contact technical support.
- **Supported configurations**—The following table identifies the supported configurations for a files and folders migration job.

Configuration	Description	Supported	Not Supported
One to one active/standby	You can migrate a single source to a single target. The target has no production activity. The source is the only server actively replicating data.	X	
One to one active/active	You cannot migrate a single source to a single target where each server acts as both a source and target actively replicating data to each other.		X
Many to one	You can migrate many source servers to one target. Replication occurs from each source to the one target. This will consolidate your source servers to a single server.	X	
One to many	You cannot migrate a single source to multiple target servers.		X
Chained	You cannot migrate a single source to a single target, where the target then acts as a source, sending the same data from the original source to a final target server.		X

Configuration	Description	Supported	Not Supported
Single server	You can migrate a single source to itself allowing data to be replicated from one location to another on the same volume or to a separate volume on the same server.	X	
Standalone to standalone	Your servers can be in a standalone to standalone configuration.	X	
Standalone to cluster	A target cluster is not a default supported configuration, but it is possible with assistance from Professional Services. If you want to migrate standalone to a target cluster, contact Sales or Professional Services. Because this is an advanced configuration, you will be referred to Professional Services for further assistance.	X	
Cluster to standalone	A source cluster is not a default supported configuration, but it is possible with assistance from Professional Services. If you want to migrate from a cluster, contact Sales or Professional Services. Because this is an advanced configuration, you will be referred to Professional Services for further assistance.	X	
Cluster to cluster	A source and target cluster is not a default supported configuration, but it is possible with assistance from Professional Services. If you want to migrate to and from clusters, contact Sales or Professional Services. Because this is an advanced configuration, you will be referred to Professional Services for further assistance.	X	

If you are using a Cluster Shared Volume (CSV), you cannot migrate any data, including a virtual machine, residing on the CSV. If you want to migrate a CSV virtual machine, you must run Carbonite Migrate from within the guest operating system of the virtual machine and create the job within the guest.

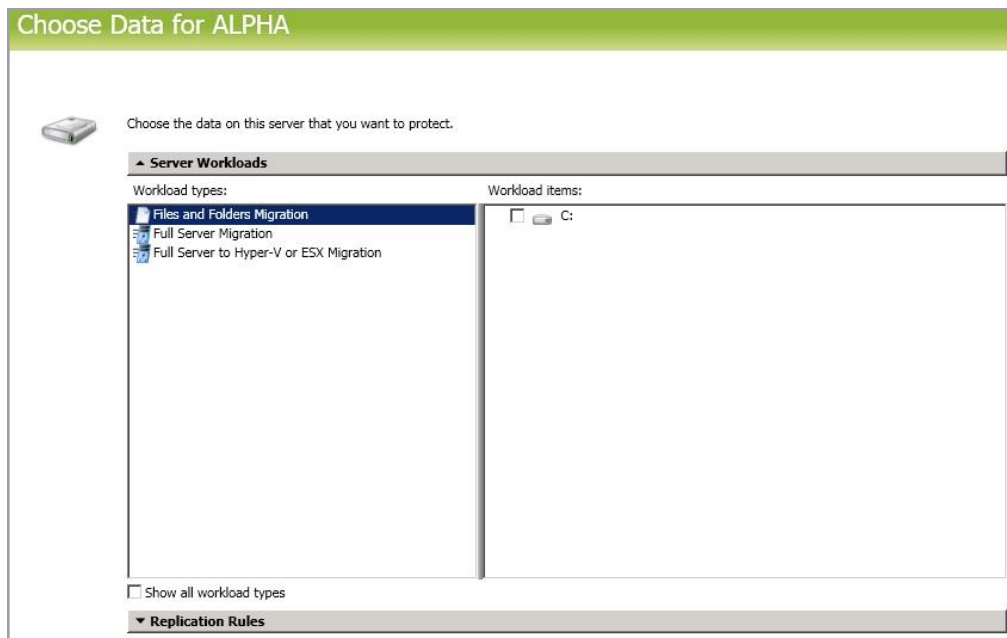
Data can be written to a target CSV.

Creating a files and folders migration job

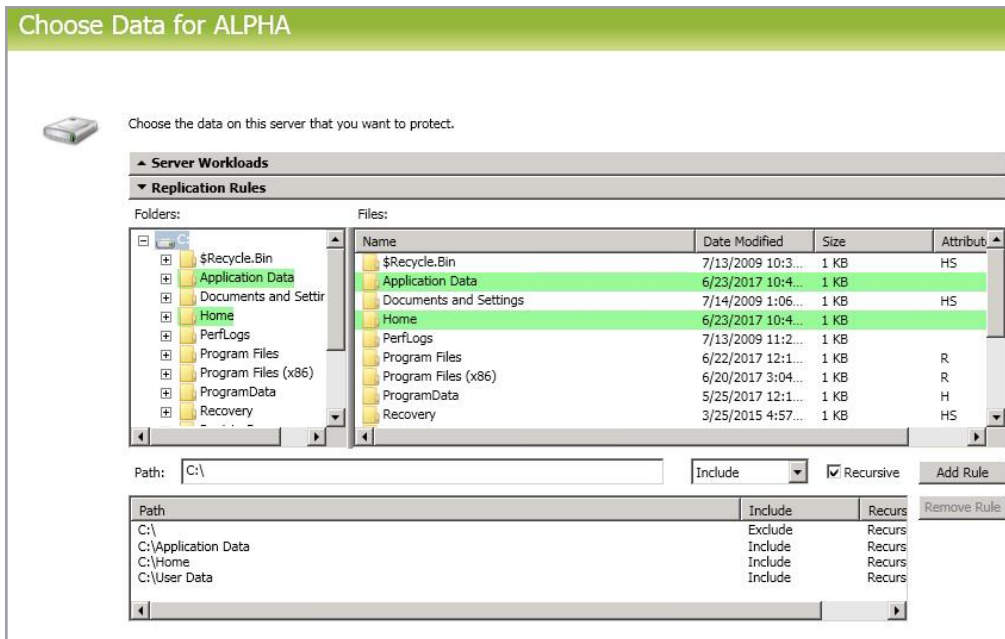
Use the following instructions to create a files and folders migration job.

1. From the **Servers** page, right-click the server you want to migrate and select **Migrate**. You can also highlight a server, click **Create a New Job** in the toolbar, then select **Migrate**.
2. Choose the type of workload that you want to migrate. Under **Server Workloads**, in the **Workload types** pane, select **Files and Folders Migration**. In the **Workload items** pane, you will see the volumes and shares (if any) for your source. Select the volumes or shares on the source that you want to migrate. You can select your files and folders in more detail in the **Replication Rules** section.

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



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



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

You can also enter wildcard rules, however you should do so carefully. Rules are applied to files that are closest in the directory tree to them. If you have rules that include multiple folders, an exclusion rule with a wild card will need to be added for each folder that it needs applied to. For example, if you want to exclude all .log files from D:\ and your rules include D:\, D:\Dir1 , and D:\Dir2, you would need to add the exclusion rule for the root and each subfolder rule. So you will need to add exclude rules for D:*.log , D:\Dir1*.log, and D:\Dir2*.log.


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



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

- Click **Next** to continue.

5. Choose your target server. This is the server that will receive the migrated data from the source.



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



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

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

6. Click **Next** to continue.



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

7. You have many options available for your files and folders migration job. Configure those options that are applicable to your environment.

Go to each page identified below to see the options available for that section of the **Set Options** page. After you have configured your options, continue with the next step on page 82.

- *General* on page 69
- *Failover Options* on page 70
- *Mirror, Verify & Orphaned Files* on page 72
- *Network Route* on page 75
- *Path Mapping* on page 76
- *Compression* on page 78
- *Bandwidth* on page 79
- *Scripts* on page 81


General

General

Job name:

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

Failover Options



- **Wait for user to initiate failover**—The cutover process can wait for you to initiate it, allowing you to control when cutover occurs. When a cutover occurs, the job will wait in the **Protecting** state for you to manually initiate the cutover process. Disable this option if you want cutover to occur immediately after the mirror is complete.
- **Shutdown source server**—Specify if you want to shut down the source server, if it is still running, before the source is cutover to the target. This option prevents identity conflicts on the network in those cases where the source and target are still both running and communicating.
- **Failover shares**—If desired, you can mirror and replicate data from your source shares to the target. During cutover, these shares will be added to the target. Keep in mind that additional steps may be needed after cutover to redirect users and/or applications to the shares on the new server. For example, you may need to rename the new server to the same name as the original server or you may need to update login scripts for the new server name.



Automatic share cutover only occurs for standard Windows file system shares. Other shares must be configured for cutover through the failover scripts or created manually on the target. See *Macintosh shares* on page 259 or *NFS Shares* on page 260 for more information.

If you are cutting over Windows shares but your source and target do not have the same drive letters, you must use the **All to One** selection under **Path Mapping** when establishing your job. Otherwise, the shares will not be created on the target during cutover.

Windows share information is automatically updated on the target every 30 seconds.

- **Target Scripts**—You can customize cutover by running scripts on the target. Scripts may contain any valid Windows command, executable, or batch file. The scripts are processed using the same account running the Double-Take Management service, unless you have identified a specific account through the server's properties. See *Script credentials* on page 51. Examples of functions specified in scripts include stopping services on the target before cutover because they may not be necessary, stopping services on the target that need to be restarted with the source's machine name and/or IP address, starting services or loading

applications that are in an idle, standby mode waiting for cutover to occur, notifying the administrator before and after cutover occurs, and so on. There are two types of cutover scripts.

- **Pre-failover script**—This script runs on the target at the beginning of the cutover process. Specify the full path and name of the script file.
- **Post-failover script**—This script runs on the target at the end of the cutover process. Specify the full path and name of the script file.
- **Arguments**—Specify a comma-separated list of valid arguments required to execute the script.
- **Delay until script completes**—Enable this option if you want to delay the cutover process until the associated script has completed. If you select this option, make sure your script handles errors, otherwise the cutover process may never complete if the process is waiting on a script that cannot complete.

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

Mirror, Verify & Orphaned Files

Mirror, Verify & Orphaned Files

Mirror Options

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

Compare file attributes. Send the attributes and bytes that differ.

Verification Options

Enable scheduled verification

Verify on this interval: 1 Days

Begin immediately

Begin at this time: 9/26/2016 10:00:43 AM

Report and comparison options

Report only

Report and mirror files

Compare file attributes and data

General Options

Calculate size of protected data upon connection

Delete orphaned files

- **Mirror Options**—Choose a comparison method and whether to mirror the entire file or only the bytes that differ in each file.
 - **Do not compare files. Send the entire file.**—Carbonite Migrate 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 option requires no time for comparison, but it can be slower because it sends the entire file. However, it is useful for configurations that have large data sets with millions of small files that are frequently changing and it is more efficient to send the entire file. You may also need to use this option if configuration management policies require sending the entire file.
 - **Compare file attributes. Send the entire file.**—Carbonite Migrate will compare file attributes and will mirror those files that have different attributes, sending the entire file. This option is the fastest comparison method, but can be slower because it sends the entire file. However, it is useful for configurations that have large data sets with millions of small files that are mostly static and not changing. You may also need to use this option if configuration management policies require sending the entire file.
 - **Compare file attributes. Send the attributes and bytes that differ.**—Carbonite Migrate will compare file attributes and will mirror only the attributes and bytes that are different. This option is the fastest comparison method and fastest mirror option. Files that have not changed can be easily skipped. Also files that are open and require a checksum mirror can be compared.
 - **Compare file attributes and data. Send the attributes and bytes that differ.**—Carbonite Migrate will compare file attributes and the file data and will mirror only the attributes and bytes that are different. This comparison method is not as fast because every file is compared, regardless of whether the file has changed or is open. However, sending only the attributes and bytes that differ is the fastest mirror 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 Carbonite Migrate 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, Carbonite Migrate 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 Carbonite Migrate 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.


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

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

If you want to move orphaned files rather than delete them, you can configure this option along with the move deleted files feature to move your orphaned files to the specified deleted files directory. See *Target server properties* on page 47 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 Carbonite Migrate 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 Carbonite Migrate log, and during difference mirrors, verifications, and restorations, all orphaned file processing messages are logged to the Carbonite Migrate log. The orphaned file log is located in the **Logging folder** specified for the source. See *Log file properties* on page 52 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



Send data to the target server using this route:
10.10.10.30

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



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

Path Mapping

Source Path	Target Path
C:\	C:\ALPHA\C' ...

Block target paths upon connection

All to One One to One

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



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


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

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

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

- **Block target paths upon connection**—You can block writing to the replica source data located on the target. This keeps the data from being changed outside of Carbonite Migrate processing. Any target paths that are blocked will be unblocked automatically during the cutover process so that users can modify data after cutover.

Compression



To help reduce the amount of bandwidth needed to transmit Carbonite Migrate 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.

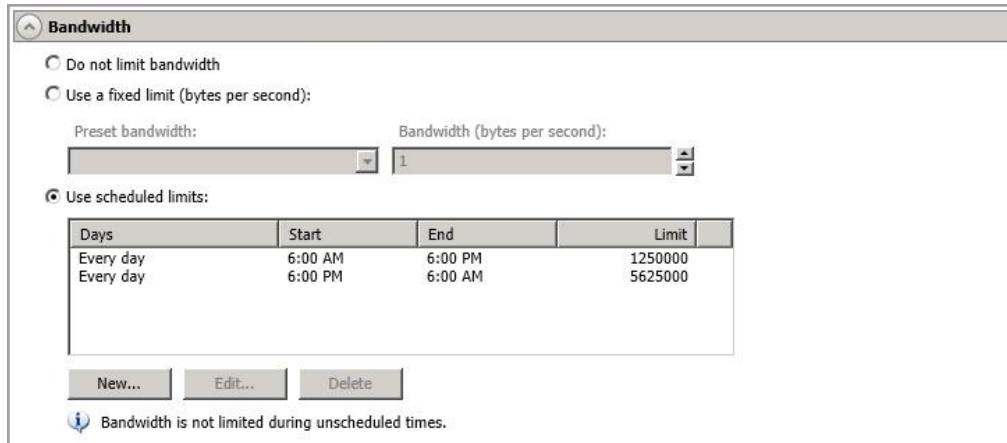
Keep in mind that the process of compressing data impacts processor usage on the source. If you notice an impact on performance while compression is enabled in your environment, either adjust to a lower level of compression, or leave compression disabled. Use the following guidelines to determine whether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



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

Bandwidth



Bandwidth


Do not limit bandwidth

Use a fixed limit (bytes per second):

Preset bandwidth: Bandwidth (bytes per second):

Use scheduled limits:

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

 Bandwidth is not limited during unscheduled times.

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



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

- **Do not limit bandwidth**—Carbonite Migrate will transmit data using 100% bandwidth availability.
- **Use a fixed limit**—Carbonite Migrate 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**—Carbonite Migrate will transmit data using a dynamic bandwidth based on the schedule you configure. Bandwidth will not be limited during unscheduled times.
 - **New**—Click **New** to create a new scheduled bandwidth limit. Specify the following information.
 - **Daytime entry**—Select this option if the start and end times of the bandwidth window occur in the same day (between 12:01 AM and midnight). The start time must occur before the end time.
 - **Overnight entry**—Select this option if the bandwidth window begins on one day and continues past midnight into the next day. The start time must be later than the end time, for example 6 PM to 6 AM.
 - **Day**—Enter the day on which the bandwidth limiting should occur. You can pick a specific day of the week, **Weekdays** to have the limiting occur Monday through Friday, **Weekends** to have the limiting occur Saturday and Sunday, or **Every day** to have the limiting repeat on all days of the week.

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



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

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

Scripts

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

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

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

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



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

8. Click **Next** to continue.
9. Carbonite Migrate 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 Carbonite Migrate correct the problem for you. For those errors that Carbonite Migrate 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 *Carbonite Availability and Carbonite Migrate Reference Guide* for details on the various Carbonite Migrate log files.

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



Jobs in a NAT environment may take longer to start.

Once a job is created, do not change the name of underlying hardware components used in the job. For example, volume and datastore names or network adapter and virtual switch names. Any component used by name in your job must continue to use that name throughout the lifetime of job. If you must change a name, you will need to delete the job and re-create it using the new component name.

Managing and controlling files and folders migration jobs

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

The jobs displayed in the right pane depend on the server group folder selected in the left pane. Every job for each server in your console session is displayed when the **Jobs on All Servers** group is selected. If you have created and populated server groups (see *Managing servers* on page 18), then only the jobs associated with the server or target servers in that server group will be displayed in the right pane.

- *Overview job information displayed in the top right pane* on page 83
- *Detailed job information displayed in the bottom right pane* on page 86
- *Job controls* on page 88


Overview job information displayed in the top right pane


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


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


Column 1 (Blank)

The first blank column indicates the state of the job.

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

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

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

 The job is in an unknown state.

Job

The name of the job

Source Server

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

Target Server

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

Job Type

Each job type has a unique job type name. This job is a Files and Folders Migration job. For a complete list of all job type names, press F1 to view the Carbonite Replication Console online help.

Activity

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

Mirror Status

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

Replication Status

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

Transmit Mode

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

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

Operating System

The job type operating system

Detailed job information displayed in the bottom right pane

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

Name

The name of the job

Target data state

- **OK**—The data on the target is in a good state.
- **Mirroring**—The target is in the middle of a mirror process. The data will not be in a good state until the mirror is complete.
- **Mirror Required**—The data on the target is not in a good state because a remirror is required. This may be caused by an incomplete or stopped mirror or an operation may have been dropped on the target.
- **Busy**—The source is low on memory causing a delay in getting the state of the data on the target.
- **Not Loaded**—Carbonite Migrate 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 38.

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 38.

Replication queue

The total number of replication bytes in the source queue

Disk queue

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

Recovery point latency

The length of time replication is behind on the target compared to the source. This is the time period of replication data that would be lost if a failure were to occur at the current time. This value represents replication data only and does not include mirroring data. If you are mirroring and failover, the data on the target will be at least as far behind as the recovery point latency. It could potentially be further behind depending on the

circumstances of the mirror. If mirroring is idle and you failover, the data will only be as far behind as the recovery point latency time.

Bytes sent

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

Bytes sent (compressed)

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

Connected since

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

Recent activity

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

Additional information

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

Job controls

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

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

View Job Details

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

Edit Job Properties

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

Delete

Stops (if running) and deletes the selected jobs.

Provide Credentials

Changes the login credentials that the job (which is on the target machine) uses to authenticate to the servers in the job. This button opens the Provide Credentials dialog box where you can specify the new account information and which servers you want to update. See *Providing server credentials* on page 31. You will remain on the **Jobs** page after updating the server credentials. If your servers use the same credentials, make sure you also update the credentials on the **Servers** page so that the Carbonite Replication Console can authenticate to the servers in the console session. See *Managing servers* on page 18.

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 Carbonite Migrate queue was not exhausted during

the time the job was paused. If the Carbonite Migrate queue was exhausted during the time the job was paused, the job will restart mirroring and replication.

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

Pause

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

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

Stop

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

Take Snapshot

Snapshots are not applicable to migration jobs.

Manage Snapshots

Snapshots are not applicable to migration jobs.

Failover or Cutover

Starts the cutover process. See *Cutting over files and folders migration jobs* on page 102 for the process and details of cutting over a files and folders migration job.

Failback

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

Restore

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

Reverse

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

Undo Failover or Cutover

Cancels a test cutover by undoing it. Undo failover does not apply to files and folders migration jobs.

View Job Log

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

Other Job Actions

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

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

When pausing a mirror, Carbonite Migrate 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, Carbonite Migrate 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.**—Carbonite Migrate 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 option requires no time for comparison, but it can be slower because it sends the entire file. However, it is useful for configurations that have large data sets with millions of small files that are frequently changing and it is more efficient to send the entire file. You may also need to use this option if configuration management policies require sending the entire file.
 - **Compare file attributes. Send the entire file.**—Carbonite Migrate will compare file attributes and will mirror those files that have different attributes, sending the entire file. This option is the fastest comparison method, but can be slower because it sends the entire file. However, it is useful for configurations that have large data sets with millions of small files that are mostly static and not changing. You may also need to use this option if configuration management policies require sending the entire file.
 - **Compare file attributes. Send the attributes and bytes that differ.**—Carbonite Migrate will compare file attributes and will mirror only the attributes and bytes that are different. This option is the fastest comparison method and fastest mirror option. Files that have not changed can be easily skipped. Also files that are open and require a checksum mirror can be compared.

- **Compare file attributes and data. Send the attributes and bytes that differ.**—Carbonite Migrate will compare file attributes and the file data and will mirror only the attributes and bytes that are different. This comparison method is not as fast because every file is compared, regardless of whether the file has changed or is open. However, sending only the attributes and bytes that differ is the fastest mirror option.
- **Calculate size of protected data before mirroring**—Specify if you want Carbonite Migrate 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**—Carbonite Migrate will transmit data using 100% bandwidth availability.
 - **Fixed bandwidth limit**—Carbonite Migrate 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 Carbonite Migrate 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 Carbonite Migrate processing, not the entire server.

While the target is paused, the Carbonite Migrate 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, Carbonite Migrate will automatically disconnect the connections and attempt to reconnect them.

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

Filter

Select a filter option from the drop-down list to only display certain jobs. You can display **Healthy jobs**, **Jobs with warnings**, or **Jobs with errors**. To clear the filter, select **All jobs**. If you have created and populated server groups, then the filter will only apply to the jobs associated with the server or target servers in that server group. See *Managing servers* on page 18.

Search

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

Overflow Chevron

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

Viewing files and folders migration job details

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

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





Job name

The name of the job

Job type

Each job type has a unique job type name. This job is a Files and Folders Migration job. For a complete list of all job type names, press F1 to view the Carbonite Replication Console online help.

Health

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

Activity

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

Connection ID

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

Transmit mode

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

Target data state

- **OK**—The data on the target is in a good state.
- **Mirroring**—The target is in the middle of a mirror process. The data will not be in a good state until the mirror is complete.
- **Mirror Required**—The data on the target is not in a good state because a remirror is required. This may be caused by an incomplete or stopped mirror or an operation may have been dropped on the target.
- **Busy**—The source is low on memory causing a delay in getting the state of the data on the target.
- **Not Loaded**—Carbonite Migrate 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 Carbonite Migrate transmissions.

Compression

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

Encryption

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

Bandwidth limit

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

Connected since

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

Additional information

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

Mirror status

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

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

Mirror percent complete

The percentage of the mirror that has been completed

Mirror remaining

The total number of mirror bytes that are remaining to be sent from the source to the target. This value may be zero if you have enabled **Mirror only changed files when source reboots** on the *Server setup properties* on page 38.

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 38.

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 Carbonite Migrate driver. Check the Event Viewer for driver related issues.
- **Unknown**—The console cannot determine the status.

Replication queue

The total number of replication bytes in the source queue

Disk queue

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

Bytes sent

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

Bytes sent compressed

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

Recovery point latency

The length of time replication is behind on the target compared to the source. This is the time period of replication data that would be lost if a failure were to occur at the current time. This value represents replication data only and does not include mirroring data. If you are mirroring and failover, the data on the target will be at least as far behind as the recovery point latency. It could potentially be further behind depending on the circumstances of the mirror. If mirroring is idle and you failover, the data will only be as far behind as the recovery point latency time.

Mirror start time

The UTC time when mirroring started

Mirror end time

The UTC time when mirroring ended

Total time for last mirror

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

Validating a files and folders migration job

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

1. From the **Jobs** page, highlight the job and click **View Job Details** in the toolbar.
2. In the **Tasks** area on the right on the **View Job Details** page, click **Validate job properties**.
3. Carbonite Migrate 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 Carbonite Migrate correct the problem for you. For those errors that Carbonite Migrate cannot correct automatically, you will need to modify the source or target to correct the error, or you can select a different target. You must revalidate the selected servers, by clicking **Recheck**, until the validation check passes without errors.

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

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

Editing a files and folders migration job

Use these instructions to edit a files and folders migration job.

1. From the **Jobs** page, highlight the job and click **Edit Job Properties** in the toolbar. (You will not be able to edit a job if you have removed the source of that job from your Carbonite Replication Console session or if you only have Carbonite Migrate monitor security access.)
2. You have the same options available for your files and folders migration job as when you created the job. See *Creating a files and folders migration job* on page 65 for details on each job option.



Changing some options may require Carbonite Migrate to automatically disconnect, reconnect, and remirror the job.

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

-
3. If you want to modify the workload items or replication rules for the job, click **Edit workload or replication rules**. Modify the **Workload item** you are protecting, if desired. Additionally, you can modify the specific **Replication Rules** for your job.

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

You can also enter wildcard rules, however you should do so carefully. Rules are applied to files that are closest in the directory tree to them. If you have rules that include multiple folders, an exclusion rule with a wild card will need to be added for each folder that it needs applied to. For example, if you want to exclude all .log files from D:\ and your rules include D:\, D:\Dir1 , and D:\Dir2, you would need to add the exclusion rule for the root and each subfolder rule. So you will need to add exclude rules for D:*.log , D:\Dir1*.log, and D:\Dir2*.log.

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

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



If you remove data from your workload and that data has already been sent to the target, you will need to manually remove that data from the target. Because the data you removed is no longer included in the replication rules, Carbonite Migrate orphan file detection cannot remove the data for you. Therefore, you have to remove it manually.

-
4. Click **Next** to continue.

5. Carbonite Migrate 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 Carbonite Migrate correct the problem for you. For those errors that Carbonite Migrate 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 *Carbonite Availability and Carbonite Migrate Reference Guide* for details on the various Carbonite Migrate log files.

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

Viewing a files and folders migration job log

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



Time	Description
6/22/2017 2:42:52 PM	Hardware IDs as follows: Source = 'bb61e75f-f6d5-4c53-9091-29017e974f2f', Target = '2d...
6/22/2017 2:42:52 PM	Completing initialization of new job 6d2bfb9-6a93-4eb4-8530-ca317cc7fcf9 (ALPHA to BE...
6/22/2017 2:42:53 PM	Initialization of job 6d2bfb9-6a93-4eb4-8530-ca317cc7fcf9 (ALPHA to BETA) complete
6/22/2017 2:42:53 PM	Changing to StoppedState from UninitializedState in response to InitializeEvent consumed...
6/22/2017 2:42:53 PM	Exited UninitializedState
6/22/2017 2:42:53 PM	Entered InitializedState
6/22/2017 2:42:53 PM	Starting monitor dfe2ee61-fde5-4afb-9ffa-17497808320c: name = FilesAndFolders_6d2bfb...
6/22/2017 2:42:53 PM	Entered StoppedState
6/22/2017 2:42:53 PM	Stopping monitor dfe2ee61-fde5-4afb-9ffa-17497808320c: Name = FilesAndFolders_6d2bfb...
6/22/2017 2:42:53 PM	Stopping share monitoring
6/22/2017 2:42:53 PM	Changing connection health to Warning
6/22/2017 2:42:53 PM	Event log entry written: '6008'.
6/22/2017 2:42:54 PM	Scheduler added new request 5b60863f-1ef2-4981-ae0d-44c0a79ead37
6/22/2017 2:42:54 PM	Deleted replication set named FilesAndFolders_6d2bfb96a934eb48530ca317cc7fcf9
6/22/2017 2:42:55 PM	Successfully created connection 92a61fa2-387a-4759-9fc8-60bc55141f08 connecting FilesA...
6/22/2017 2:42:55 PM	Attaching to engine connection on 172.31.206.200:6325 with following criteria:Guid = '92a...
6/22/2017 2:42:55 PM	Waiting 00:10:00 for source endpoint of 'FilesAndFolders_6d2bfb96a934eb48530ca317cc...
6/22/2017 2:42:55 PM	Established source endpoint of '172.31.206.200:6320' for engine connection with replicatio...
6/22/2017 2:42:55 PM	Updating failover options
6/22/2017 2:42:58 PM	The Double-Take engine is initialized.
6/22/2017 2:42:58 PM	Double-Take is NOT licensed to monitor or assume the identity of another machine.
6/22/2017 2:42:58 PM	The Double-Take engine source module is initialized.
6/22/2017 2:42:58 PM	The Double-Take engine target module is initialized.
6/22/2017 2:43:01 PM	Updating IPAddresses (Request - 5b60863f-1ef2-4981-ae0d-44c0a79ead37, WorkflowId - ...
6/22/2017 2:43:01 PM	Starting share monitoring
6/22/2017 2:43:03 PM	Changing targetActivationCode health to Ok
6/22/2017 2:43:03 PM	Event log entry written: '6004'.
6/22/2017 2:43:03 PM	Changing to ConnectedState from StoppedState in response to StartSucceededEvent consu...
6/22/2017 2:43:03 PM	Exited StoppedState
6/22/2017 2:43:03 PM	Entered ConnectedState
6/22/2017 2:43:03 PM	Subscribing to engine connection.
6/22/2017 2:43:03 PM	Changing sourceActivationCode health to Ok
6/22/2017 2:43:03 PM	Changing connection health to Ok
6/22/2017 2:43:03 PM	Changing to SynchronizedState from ConnectedState in response to MirrorCompletedEvent...
6/22/2017 2:43:03 PM	Entered ProtectingState
6/22/2017 2:43:03 PM	Starting monitor dfe2ee61-fde5-4afb-9ffa-17497808320c: name = FilesAndFolders_6d2bfb...
6/22/2017 2:43:03 PM	Entered SynchronizedState
6/22/2017 2:43:03 PM	Persisting shares
6/22/2017 2:43:03 PM	Event log entry written: '6008'.

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.

Cutting over files and folders migration jobs

When the migration mirror has completed, the migration job may or may not terminate automatically depending on your selection for **Wait for user intervention before cutover**. If you disabled user intervention, the migration job will automatically terminate to complete the migration process. If you enabled user intervention, when the migration mirror is complete, the status will change to **Protecting**. Use this time to complete any necessary tasks. When you are ready to complete the migration, use the following instructions to cutover.

1. On the **Jobs** page, highlight the job that you want to cutover and click **Failover, Cutover, or Recover** in the toolbar.
2. Select the type of cutover to perform.
 - **Cutover live data**—Select this option to initiate a full, live cutover using the current data on the target. The source may be automatically shut down if it is still running, depending on your job configuration.
 - **Perform test cutover**—This option is not applicable to files and folders migration jobs.
 - **Cutover to a snapshot**—This option is not available for migration jobs.
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 cutover begins. The advantage to this option is that all of the data that the target has received will be applied before cutover 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 cutover will begin immediately. The advantage to this option is that cutover will occur immediately. The disadvantage is that any data in the target queue will be lost.
4. When you are ready to begin cutover, click **Cutover**.



If your source was disconnected from the network during the cutover process, the original job that was mirroring the data will still be active on that machine. Do not bring that machine back on the network. If you do, the original job will attempt to begin mirroring data again which could result in data loss. While the machine is still disconnected from the network, stop the Double-Take service, delete the file connect.sts located in the Carbonite Migrate installation directory, restart the Double-Take service, and bring the machine on the network if desired.

If you selected share cutover during your files and folders migration job creation, Carbonite Migrate will mirror and replicate data from your source shares to the target. During cutover, these shares will be added to the target. Keep in mind that additional steps may be needed after cutover to redirect users and/or applications to the shares on the new server. For example, you may need to rename the new server to the same name as the original server or you may need to update login scripts for the new server name.

Because Windows 64-bit has a strict driver signing policy, if you get a stop code 0x7b after failover, you may have drivers failing to load because the driver signatures are failing the policy. In this case, reboot the server and press F8. Choose the option to not enforce the

driver signing policy. If this allows the system to boot, then the problem is being caused by the cat file signature mismatch. If your system still fails to boot, contact technical support.

Chapter 6 Full server migration

Create a full server migration job when you want to migrate the entire source, including the server's system state and applications.

- *Full server migration requirements* on page 105—Full server migration includes specific requirements for this type of migration.
- *Creating a full server migration job* on page 112—This section includes step-by-step instructions for creating a full server migration job.
- *Managing and controlling full server migration jobs* on page 131—You can view status information about your full server migration job.
- *Cutting over full server migration jobs* on page 150—Use this section when you are ready to cutover from your source to your target, which will become your new source.

Full server migration requirements

Use these requirements below for full server migration. Keep in mind that a target server may meet these requirements but may not be suitable to stand-in for a source after cutover. See *Target compatibility* on page 110 for additional information regarding an appropriate target server for your particular source.

- **Operating system**—The following operating systems are supported for full server jobs.
 - Windows 2016 and Server Core 2016
 - Windows 2012 R2 and Server Core 2012 R2
 - Windows 2012 and Server Core 2012
 - Windows 2008 R2 Service Pack 1 or later and Server Core 2008 R2 Service Pack 1 or later



Windows 2016 support is for the primary operating system features available in Windows 2012. New operating system features specific to Windows 2016, such as Nano Server, Windows Containers, and so on, are not supported.

Due to VMware limitations, Windows 2016 is not supported on ESX 5.5.

Server Core operating systems are only supported in a matching Server Core to Server Core configuration.

DNS updates are not supported for Server Core servers.

Full server migration of a Hyper-V server is not supported.

If you are using Windows Storage Server Edition, you will need to check with your NAS vendor to verify if there are technical or license restrictions on cutting over an image of a server to different hardware.

-
- **Operating system language**—Your servers must be running the same Windows localized version. For example, if your source is running an English language version of Windows, your target must also be running an English language version of Windows. If your source is running a Japanese language version of Windows, your target must also be running a Japanese language version of Windows. This applies to all localized languages.
 - **Cloud**—Carbonite Migrate can be used to migrate a server to an existing server in the cloud. Keep in mind that you should enable appropriate security measures, like VPN, to protect your data as it migrates to the cloud.
 - **Boot mode**—The boot mode can be EFI/UEFI or BIOS.
 - **Source and target preparation**—Uninstall any applications or operating system features that are not needed from both your source and target. For example, unused language packs will slow down cutover since there are thousands of extra files that need to be examined. Ideally, your target should be as clean and simple a configuration as possible.



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

- **File system**—Carbonite Migrate supports the NTFS file system. On Windows 2016 servers, ReFS is also supported. FAT and FAT32 are not supported. For detailed information on other file system capabilities, see *Mirroring and replication capabilities* on page 7.
 - **Microsoft .NET Framework**—Microsoft .NET Framework version 4.5.1 or later is required on the source and target.
 - **System memory**—The minimum system memory on each server is 1 GB.
 - **Disk space for program files**—This is the amount of disk space needed for the Carbonite Migrate program files. The amount depends on your operating system version and ranges from 350-500 MB.
-



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

Make sure you have additional disk space for Carbonite Migrate queuing, logging, and so on.

- **Server name**—Carbonite Migrate includes Unicode file system support, but your server name must still be in ASCII format. Additionally, all Carbonite Migrate servers and appliances must have a unique server name.
- **Time**—The clock on your Carbonite Migrate servers must be within a few minutes of each other, relative to UTC. Large time skews (more than five minutes) will cause Carbonite Migrate errors.
- **Protocols and networking**—Your servers must meet the following protocol and networking requirements.
 - Your servers must have TCP/IP with static IP addressing.
 - IPv4 only configurations are supported, IPv4 and IPv6 are supported in combination, however IPv6 only configurations are not supported
 - If your target server only supports DHCP, for example Windows Azure, keep in mind the following caveats.
 - A target reboot may or may not cause a job error depending on if a new address is assigned by DHCP.
 - Do not disable the DHCP client service on the source. Otherwise, when cutover occurs the DHCP client will not start and an IP address cannot be assigned.
 - If you are using IPv6 on your servers, your clients must be run from an IPv6 capable machine.
 - In order to properly resolve IPv6 addresses to a hostname, a reverse lookup entry should be made in DNS.

- If you are using Carbonite Migrate over a WAN and do not have DNS name resolution, you will need to add the host names to the local hosts file on each server running Carbonite Migrate.
- **Network adapters**—Microsoft NIC teaming for Windows 2012 is supported, however no third party NIC teaming solutions are supported.
- **NAT support**—Carbonite Migrate supports IP and port forwarding in NAT environments with the following caveats.
 - Only IPv4 is supported.
 - Only standalone servers are supported.
 - Make sure you have added your server to the Carbonite Replication Console using the correct public or private IP address. The name or IP address you use to add a server to the console is dependent on where you are running the console. Specify the private IP address of any servers on the same side of the router as the console. Specify the public IP address of any servers on the other side of the router as the console.
 - DNS failover and updates will depend on your configuration
 - Only the source or target can be behind a router, not both.
 - The DNS server must be routable from the target
- **Reverse lookup zone**—If you are using a DNS reverse lookup zone, then it must be Active Directory integrated. Carbonite Migrate is unable to determine if this integration exists and therefore cannot warn you during job creation if it doesn't exist.
- **DNS updates**—You can 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. Carbonite Migrate 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.

- **Windows firewall**—If you have Windows firewall enabled on your servers, there are two requirements for the Windows firewall configuration.
 - The Carbonite Migrate installation program will automatically attempt to configure ports 6320, 6325, and 6326 for Carbonite Migrate. If you cancel this step, you will have to configure the ports manually.
 - If you are using the Carbonite Replication Console to push installations out to your Windows servers, you will have to open firewall ports for WMI (Windows Management

Instrumentation), which uses RPC (Remote Procedure Call). By default, RPC will use ports at random above 1024, and these ports must be open on your firewall. RPC ports can be configured to a specific range by specific registry changes and a reboot. See the [Microsoft Knowledge Base article 154596](#) for instructions. Additionally, you will need to open firewall ports for SMB (server message block) communications which uses ports 135-139 and port 445, and you will need to open File and Printer Sharing. As an alternative, you can disable the Windows firewall temporarily until the push installations are complete.

See *Firewalls* on page 255 for instructions on handling firewalls in your environment.

- **Windows Management Instrumentation (WMI)**—Carbonite Migrate is dependent on the WMI service. If you do not use this service in your environment, contact technical support.
- **Supported configurations**—The following table identifies the supported configurations for a full server migration job.

Configuration	Description	Supported	Not Supported
One to one active/standby	You can migrate a single source to a single target. The target has no production activity. The source is the only server actively replicating data.	X	
One to one active/active	You cannot migrate a single source to a single target where each server acts as both a source and target actively replicating data to each other.		X
Many to one	You cannot migrate many source servers to one target server.		X
One to many	You cannot migrate a single source to multiple target servers.		X
Chained	You cannot migrate a single source to a single target, where the target then acts as a source, sending the same data from the original source to a final target server.		X
Single server	You cannot migrate a single source to itself.		X
Standalone to standalone	Your servers can be in a standalone to standalone configuration.	X	
Standalone to cluster	A target cluster is not a default supported configuration, but it is possible with assistance from Professional Services. If you want to migrate standalone to a target cluster, contact Sales or Professional Services. Because this is an advanced configuration, you will be referred to Professional Services for further assistance.	X	

Configuration	Description	Supported	Not Supported
Cluster to standalone	A source cluster is not a default supported configuration, but it is possible with assistance from Professional Services. If you want to migrate from a cluster, contact Sales or Professional Services. Because this is an advanced configuration, you will be referred to Professional Services for further assistance.	X	
Cluster to cluster	A source and target cluster is not a default supported configuration, but it is possible with assistance from Professional Services. If you want to migrate to and from clusters, contact Sales or Professional Services. Because this is an advanced configuration, you will be referred to Professional Services for further assistance.	X	

If you are using a Cluster Shared Volume (CSV), you cannot migrate any data, including a virtual machine, residing on the CSV. If you want to migrate a CSV virtual machine, you must run Carbonite Migrate from within the guest operating system of the virtual machine and create the job within the guest.

Data can be written to a target CSV.

Target compatibility

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



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

- **Storage Server Edition**—If you are using Windows Storage Server Edition, you will need to check with your NAS vendor to verify if there are technical or license restrictions on cutting over an image of a server to different hardware.
- **Windows Azure**—Because Windows Azure uses remote desktop (RDP) for console connections to virtual machines running on Azure, if your target is running on Windows Azure, you must have remote desktop enabled on the source or the target will be inaccessible after cutover.
- **Server role**—The target cannot be a domain controller. Ideally, the target should not host any functionality (file server, application server, and so on) because the functionality will be removed when cutover occurs.

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

- **Processors**—There are no limits on the number or speed of the processors, but the source and the target should have at least the same number of processors. If the target has fewer processors or slower speeds than the source, there will be performance impacts for the users after cutover.
- **Memory**—The target memory should be within 25% (plus or minus) of the source. If the target has much less memory than the source, there will be performance impacts for the users after cutover.
- **Network adapters**—You must map at least one NIC from the source to one NIC on the target. If you have NICs on the source that are not being used, it is best to disable them. If the source has

more NICs than the target, some of the source NICs will not be mapped to the target. Therefore, the IP addresses associated with those NICs will not be available after cutover. If there are more NICs on the target than the source, the additional NICs will still be available after cutover and will retain their pre-cutover network settings.

- **File system format**—The source and the target must have the same NTFS or ReFS file system format on each server. FAT and FAT32 are no longer supported.
- **Logical volumes**—There are no limits to the number of logical volumes, although you are bound by operating system limits. For each volume you are migrating on the source, the target must have a matching volume. For example, if you are migrating drives C: and D: on the source, the target cannot have drives D: and E:. In this case, the target must also have drives C: and D:. Additional target volumes are preserved and available after cutover with all data still accessible.
- **System path**—The source and the target must have the same system path. The system path includes the location of the Windows files, Program Files, and Documents and Settings.
- **Disk space**—The target must have enough space to store the data from the source. This amount of disk space will depend on the applications and data files you are protecting. The more data you are protecting, the more disk space you will need. The target must also have enough space to store, process, and apply the source's system state data. The size of the system state will depend on the operating system and architecture.

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

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

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

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

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

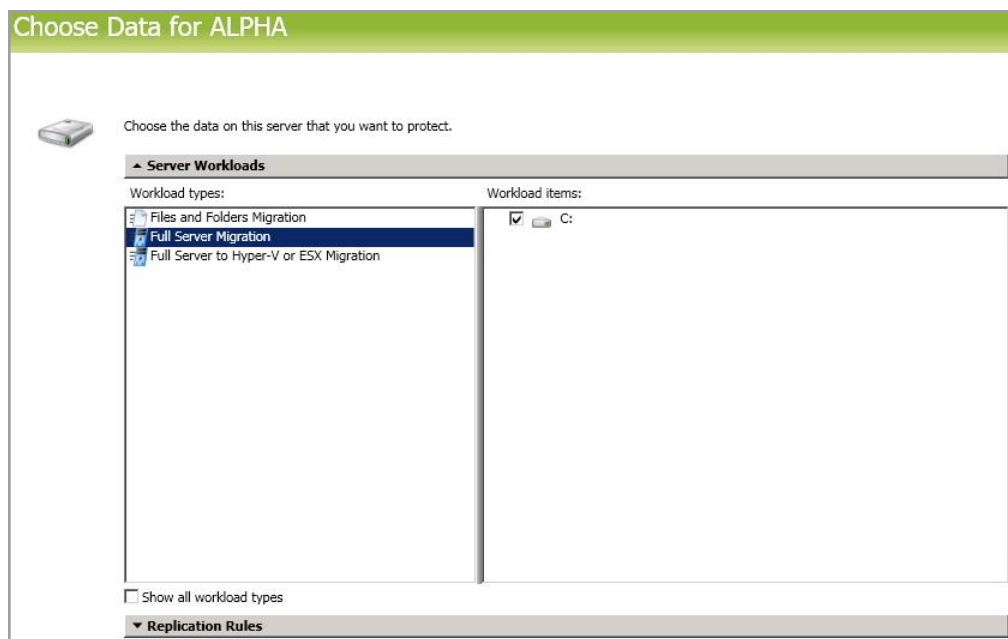
These minimums are for a clean operating system installation. Operating system customizations, installed applications, and user data will increase the disk space requirement.

Creating a full server migration job

Use the following instructions to create a full server migration job.

1. From the **Servers** page, right-click the server you want to migrate and select **Migrate**. You can also highlight a server, click **Create a New Job** in the toolbar, then select **Migrate**.
2. Choose the type of workload that you want to migrate. Under **Server Workloads**, in the **Workload types** pane, select **Full Server Migration**. In the **Workload items** pane, select the volumes on the source that you want to migrate.

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



3. By default, Carbonite Migrate selects your entire source for migration. If desired, click the **Replication Rules** heading and expand the volumes under **Folders**. You will see that Carbonite Migrate automatically excludes particular files that cannot be used during the migration. If desired, you can exclude other files that you do not want to migrate, but be careful when excluding data. Excluded volumes, folders, and/or files may compromise the integrity of your installed applications. There are some volumes, folders, and files (identified in *italics text*) that you will be unable to exclude, because they are required for migration. For example, the boot files cannot be excluded because that is where the system state information is stored.

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

You can also enter wildcard rules, however you should do so carefully. Rules are applied to files that are closest in the directory tree to them. If you have rules that include multiple folders, an exclusion rule with a wild card will need to be added for each folder that it needs applied to. For example, if you want to exclude all .log files from D:\ and your rules include D:\, D:\Dir1 , and D:\Dir2, you would need to add the exclusion rule for the root and each subfolder rule. So you will need to add exclude rules for D:*.log , D:\Dir1*.log, and D:\Dir2*.log.


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



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

If IIS is being used as a hardware platform manager by your hardware vendor on both the source and target, you need to remove the INetPub directory from replication under the **Replication Rules** heading. If IIS is being used as a software application on your source but as a hardware platform manager by your hardware vendor on your target, you need to add the INetPub directory to the **Staged Folders Options** list on the **Set Options** page later in this workflow.

4. Click **Next** to continue.
5. Choose your target server. This is the server that, after the migration, will become your new source.



- **Current Servers**—This list contains the servers currently available in your console session. Servers that are not licensed for the workflow you have selected and those not

applicable to the workload type you have selected will be filtered out of the list. Select your target server from the list. If the server you are looking for is not displayed, enable **Show all servers**. The servers in red are not available for the source server or workload type you have selected. Hover your mouse over an unavailable server to see a reason why this server is unavailable.

- **Find a New Server**—If the server you need is not in the **Current Servers** list, click the **Find a New Server** heading. From here, you can specify a server along with credentials for logging in to the server. If necessary, you can click **Browse** to select a server from a network drill-down list.



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

When specifying credentials for a new server, specify a user that is a member of the local Double-Take Admin and local administrator security groups. If your source is the only domain controller in your network, the account must also be a local account in the local administrators group on the target. If you want Carbonite Migrate to update DNS during cutover, the account must be a member of the Domain Admins group. If your security policies do not allow use of this group, see the *Special network configurations* chapter of the *Carbonite Availability User's Guide*, in the DNS section, and use the instructions under the Carbonite Availability DFO utility to use a non-Domain Admins account.

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



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

-
7. You have many options available for your full server migration job. Configure those options that are applicable to your environment.

Go to each page identified below to see the options available for that section of the **Set Options** page. After you have configured your options, continue with the next step on page 130.

- *General* on page 115
- *Failover Options* on page 116
- *Failover Identity* on page 117
- *Network Adapter Options* on page 119
- *Mirror, Verify & Orphaned Files* on page 120
- *Network Route* on page 123
- *Staging Folder Options* on page 124
- *Target Services* on page 125
- *Compression* on page 126
- *Bandwidth* on page 127
- *Scripts* on page 129

General



The screenshot shows a window titled "General" with a small upward-pointing arrow icon on the left. Below the title bar, there is a label "Job name:" followed by a text input field. The input field contains the text "alpha to beta".

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

Failover Options

Failover Options

- Wait for user to initiate failover
- Shutdown source server
- Change target ports to match source during cutover

Target scripts

Pre-failover script: ... Arguments:

Delay failover until script completes

Post-failover script: ... Arguments:

- **Wait for user to initiate failover**—The cutover process can wait for you to initiate it, allowing you to control when cutover occurs. When a cutover occurs, the job will wait in the **Protecting** state for you to manually initiate the cutover process. Disable this option if you want cutover to occur immediately after the mirror is complete.
- **Shutdown source server**—Specify if you want to shut down the source server, if it is still running, before the source is cutover to the target. This option prevents identity conflicts on the network in those cases where the source and target are still both running and communicating.
- **Change target ports to match source during failover**—This option allows you to specify if you want the Carbonite Migrate ports on the target to be updated to match the source during cutover. This option is useful in NAT environments where the public port does not match the private port.
- **Target Scripts**—You can customize cutover by running scripts on the target. Scripts may contain any valid Windows command, executable, or batch file. The scripts are processed using the same account running the Double-Take Management service, unless you have identified a specific account through the server's properties. See *Script credentials* on page 51. Examples of functions specified in scripts include stopping services on the target before cutover because they may not be necessary, stopping services on the target that need to be restarted with the source's machine name and/or IP address, starting services or loading applications that are in an idle, standby mode waiting for cutover to occur, notifying the administrator before and after cutover occurs, and so on. There are two types of cutover scripts.
 - **Pre-failover script**—This script runs on the target at the beginning of the cutover process. Specify the full path and name of the script file.
 - **Post-failover script**—This script runs on the target at the end of the cutover process. Specify the full path and name of the script file.
 - **Arguments**—Specify a comma-separated list of valid arguments required to execute the script.
 - **Delay until script completes**—Enable this option if you want to delay the cutover process until the associated script has completed. If you select this option, make sure your script handles errors, otherwise the cutover process may never complete if the process is waiting on a script that cannot complete.

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

Failover Identity

Failover Identity

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

Retain target network configuration (Recommended for WAN configurations)

Update DNS server

DNS Options

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

These DNS servers will be updated during failover:

112.42.48.9 Remove

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

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

Update TTL (seconds):
300

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



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

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

- **Update DNS server**—Specify if you want Carbonite Migrate to update your DNS server on cutover. If DNS updates are made, the DNS records will be locked during cutover. Be sure and review the job requirements for updating DNS.
-



DNS updates are not available for Server Core servers.

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.

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


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



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

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.

Network Adapter Options



The screenshot shows a dialog box titled "Network Adapter Options" with a sub-header "Map source network adapters to target network adapters:". Below this is a table with two columns: "Source Network Adapter" and "Target Network Adapter".

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

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

Mirror, Verify & Orphaned Files

The screenshot shows the 'Mirror, Verify & Orphaned Files' configuration window. It is divided into three sections: Mirror Options, Verification Options, and General Options. The Mirror Options section has a dropdown menu set to 'Compare file attributes. Send the attributes and bytes that differ.' The Verification Options section includes a checkbox for 'Enable scheduled verification', a 'Verify on this interval' field set to '1' with a 'Days' dropdown, and radio buttons for 'Begin immediately', 'Begin at this time' (set to 3/23/2017 10:53:32 AM), and 'Report and comparison options' (with 'Report only' and 'Report and mirror files' options). A dropdown menu at the bottom of this section is set to 'Compare file attributes and data'. The General Options section has checkboxes for 'Calculate size of protected data upon connection' and 'Delete orphaned files', both of which are checked.

- **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.**—Carbonite Migrate 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 option requires no time for comparison, but it can be slower because it sends the entire file. However, it is useful for configurations that have large data sets with millions of small files that are frequently changing and it is more efficient to send the entire file. You may also need to use this option if configuration management policies require sending the entire file.
 - **Compare file attributes. Send the attributes and bytes that differ.**—Carbonite Migrate will compare file attributes and will mirror only the attributes and bytes that are different. This option is the fastest comparison method and fastest mirror option. Files that have not changed can be easily skipped. Also files that are open and require a checksum mirror can be compared.
 - **Compare file attributes and data. Send the attributes and bytes that differ.**—Carbonite Migrate will compare file attributes and the file data and will mirror only the attributes and bytes that are different. This comparison method is not as fast because every file is compared, regardless of whether the file has changed or is open. However, sending only the attributes and bytes that differ is the fastest mirror 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 Carbonite Migrate



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, Carbonite Migrate 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 Carbonite Migrate 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.

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

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

If you want to move orphaned files rather than delete them, you can configure this option along with the move deleted files feature to move your orphaned files to the specified deleted files directory. See *Target server properties* on page 47 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 Carbonite Migrate 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 Carbonite Migrate log, and during difference mirrors, verifications, and restorations, all orphaned file processing messages are logged to the Carbonite Migrate log. The orphaned file log is located in the **Logging folder** specified for the source. See *Log file properties* on page 52 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



The screenshot shows a window titled "Network Route" with a small upward-pointing arrow icon on the left. Below the title bar, the text "Send data to the target server using this route:" is displayed. Underneath this text is a dropdown menu with a small downward-pointing arrow on the right side. The dropdown menu is currently open, showing the IP address "10.10.10.30".

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

Staging Folder Options




- **Select additional folders from the source that need to be staged**—Applications running on the target that cannot be stopped will cause retry operations because Carbonite Migrate will be unable to write to open application files. In this case, you will want to mirror those application files to a staging location instead of their actual location. Generally, this will only apply to applications that are not installed in the Windows Program Files directory. In this case, click **Add** and specify the folder that you want staged. Any staged folders will be applied to their actual installation location during cutover.



If IIS is being used as a software application on your source but as a hardware platform manager by your hardware vendor on your target, you need to add the INetPub directory to the **Staged Folders Options** list. If IIS is being used as a hardware platform manager by your hardware vendor on both the source and target, you need to go to the **Choose Data** page and remove the INetPub directory from replication under the **Replication Rules** heading.

- **Show system state and profile folders**—This option displays the list of essential system state and profile folders that will be staged automatically. These essential items are displayed in a lighter color than folders you have manually added, and they cannot be removed from the list.

Target Services




- **Services to leave running on the target server during protection**—Carbonite Migrate controls which services are running and stopped on the target during protection. You can specify which services you want to keep running by clicking **Add** and selecting a service from the list. If you want to remove a service from the list, highlight it and click **Remove**.



Services are stopped on the target to protect against retry operations. Do not leave services running unless absolutely necessary.

-
- **Show essential services**—This option displays the list of essential services that will remain running on the target. The essential services are displayed in a lighter color than services you have manually added. The essential services cannot be removed from the list.

Compression



To help reduce the amount of bandwidth needed to transmit Carbonite Migrate 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 Carbonite Migrate data.



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

Bandwidth



Bandwidth


Do not limit bandwidth

Use a fixed limit (bytes per second):

Preset bandwidth: Bandwidth (bytes per second):

Use scheduled limits:

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

 Bandwidth is not limited during unscheduled times.

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



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

- **Do not limit bandwidth**—Carbonite Migrate will transmit data using 100% bandwidth availability.
- **Use a fixed limit**—Carbonite Migrate 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**—Carbonite Migrate will transmit data using a dynamic bandwidth based on the schedule you configure. Bandwidth will not be limited during unscheduled times.
 - **New**—Click **New** to create a new scheduled bandwidth limit. Specify the following information.
 - **Daytime entry**—Select this option if the start and end times of the bandwidth window occur in the same day (between 12:01 AM and midnight). The start time must occur before the end time.
 - **Overnight entry**—Select this option if the bandwidth window begins on one day and continues past midnight into the next day. The start time must be later than the end time, for example 6 PM to 6 AM.
 - **Day**—Enter the day on which the bandwidth limiting should occur. You can pick a specific day of the week, **Weekdays** to have the limiting occur Monday through Friday, **Weekends** to have the limiting occur Saturday and Sunday, or **Every day** to have the limiting repeat on all days of the week.

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



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

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

Scripts

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

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

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

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



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

8. Click **Next** to continue.
9. Carbonite Migrate 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 Carbonite Migrate correct the problem for you. For those errors that Carbonite Migrate 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 *Carbonite Availability and Carbonite Migrate Reference Guide* for details on the various Carbonite Migrate log files.

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



Jobs in a NAT environment may take longer to start.

Once a job is created, do not change the name of underlying hardware components used in the job. For example, volume and datastore names or network adapter and virtual switch names. Any component used by name in your job must continue to use that name throughout the lifetime of job. If you must change a name, you will need to delete the job and re-create it using the new component name.

Managing and controlling full server migration jobs

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

The jobs displayed in the right pane depend on the server group folder selected in the left pane. Every job for each server in your console session is displayed when the **Jobs on All Servers** group is selected. If you have created and populated server groups (see *Managing servers* on page 18), then only the jobs associated with the server or target servers in that server group will be displayed in the right pane.

- *Overview job information displayed in the top right pane* on page 131
- *Detailed job information displayed in the bottom right pane* on page 134
- *Job controls* on page 136

Overview job information displayed in the top right pane

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

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

Column 1 (Blank)

The first blank column indicates the state of the job.



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



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



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



The job is in an unknown state.

Job

The name of the job

Source Server

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

Target Server

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

Job Type

Each job type has a unique job type name. This job is a Full Server Migration job. For a complete list of all job type names, press F1 to view the Carbonite Replication Console online help.

Activity

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

Mirror Status

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

Replication Status

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

Transmit Mode

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

Operating System

The job type operating system

Detailed job information displayed in the bottom right pane

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

Name

The name of the job

Target data state

- **OK**—The data on the target is in a good state.
- **Mirroring**—The target is in the middle of a mirror process. The data will not be in a good state until the mirror is complete.
- **Mirror Required**—The data on the target is not in a good state because a remirror is required. This may be caused by an incomplete or stopped mirror or an operation may have been dropped on the target.
- **Busy**—The source is low on memory causing a delay in getting the state of the data on the target.
- **Not Loaded**—Carbonite Migrate 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 38.

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 38.

Replication queue

The total number of replication bytes in the source queue

Disk queue

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

Recovery point latency

The length of time replication is behind on the target compared to the source. This is the time period of replication data that would be lost if a failure were to occur at the current time. This value represents replication data only and does not include mirroring data. If you are mirroring and failover, the data on the target will be at least as far behind as the recovery point latency. It could potentially be further behind depending on the

circumstances of the mirror. If mirroring is idle and you failover, the data will only be as far behind as the recovery point latency time.

Bytes sent

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

Bytes sent (compressed)

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

Connected since

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

Recent activity

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

Additional information

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

Job controls

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

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

View Job Details

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

Edit Job Properties

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

Delete

Stops (if running) and deletes the selected jobs.

Provide Credentials

Changes the login credentials that the job (which is on the target machine) uses to authenticate to the servers in the job. This button opens the Provide Credentials dialog box where you can specify the new account information and which servers you want to update. See *Providing server credentials* on page 31. You will remain on the **Jobs** page after updating the server credentials. If your servers use the same credentials, make sure you also update the credentials on the **Servers** page so that the Carbonite Replication Console can authenticate to the servers in the console session. See *Managing servers* on page 18.

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 Carbonite Migrate queue was not exhausted during

the time the job was paused. If the Carbonite Migrate queue was exhausted during the time the job was paused, the job will restart mirroring and replication.

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

Pause

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

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

Stop

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

Take Snapshot

Snapshots are not applicable to migration jobs.

Manage Snapshots

Snapshots are not applicable to migration jobs.

Failover or Cutover

Starts the cutover process. See *Cutting over full server migration jobs* on page 150 for the process and details of cutting over a full server migration job.

Failback

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

Restore

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

Reverse

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

Undo Failover or Cutover

Cancels a test cutover by undoing it. Undo failover does not apply to full server migration jobs.

View Job Log

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

Other Job Actions

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

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

When pausing a mirror, Carbonite Migrate 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, Carbonite Migrate 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.**—Carbonite Migrate 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 option requires no time for comparison, but it can be slower because it sends the entire file. However, it is useful for configurations that have large data sets with millions of small files that are frequently changing and it is more efficient to send the entire file. You may also need to use this option if configuration management policies require sending the entire file.
 - **Compare file attributes. Send the attributes and bytes that differ.**—Carbonite Migrate will compare file attributes and will mirror only the attributes and bytes that are different. This option is the fastest comparison method and fastest mirror option. Files that have not changed can be easily skipped. Also files that are open and require a checksum mirror can be compared.
 - **Compare file attributes and data. Send the attributes and bytes that differ.**—Carbonite Migrate will compare file attributes and the file data and will mirror only the attributes and bytes that are different. This comparison method is not as fast because every file is compared, regardless of whether the file has changed or is open. However, sending only the attributes and bytes that differ is the fastest mirror option.
- **Calculate size of protected data before mirroring**—Specify if you want

Carbonite Migrate 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**—Carbonite Migrate will transmit data using 100% bandwidth availability.
 - **Fixed bandwidth limit**—Carbonite Migrate 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 Carbonite Migrate 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 Carbonite Migrate processing, not the entire server.

While the target is paused, the Carbonite Migrate 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, Carbonite Migrate will automatically disconnect the connections and attempt to reconnect them.

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

Filter

Select a filter option from the drop-down list to only display certain jobs. You can display **Healthy jobs**, **Jobs with warnings**, or **Jobs with errors**. To clear the filter, select **All jobs**. If you have created and populated server groups, then the filter will only apply to the jobs associated with the server or target servers in that server group. See *Managing servers* on page 18.

Search

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

Overflow Chevron

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

Viewing full server migration job details

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

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





Job name

The name of the job

Job type

Each job type has a unique job type name. This job is a Full Server Migration job. For a complete list of all job type names, press F1 to view the Carbonite Replication Console online help.

Health

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

Activity

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

Connection ID

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

Transmit mode

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

Target data state

- **OK**—The data on the target is in a good state.
- **Mirroring**—The target is in the middle of a mirror process. The data will not be in a good state until the mirror is complete.
- **Mirror Required**—The data on the target is not in a good state because a remirror is required. This may be caused by an incomplete or stopped mirror or an operation may have been dropped on the target.
- **Busy**—The source is low on memory causing a delay in getting the state of the data on the target.
- **Not Loaded**—Carbonite Migrate 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 Carbonite Migrate transmissions.

Compression

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

Encryption

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

Bandwidth limit

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

Connected since

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

Additional information

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

Mirror status

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

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

Mirror percent complete

The percentage of the mirror that has been completed

Mirror remaining

The total number of mirror bytes that are remaining to be sent from the source to the target. This value may be zero if you have enabled **Mirror only changed files when source reboots** on the *Server setup properties* on page 38.

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 38.

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 Carbonite Migrate driver. Check the Event Viewer for driver related issues.
- **Unknown**—The console cannot determine the status.

Replication queue

The total number of replication bytes in the source queue

Disk queue

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

Bytes sent

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

Bytes sent compressed

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

Recovery point latency

The length of time replication is behind on the target compared to the source. This is the time period of replication data that would be lost if a failure were to occur at the current time. This value represents replication data only and does not include mirroring data. If you are mirroring and failover, the data on the target will be at least as far behind as the recovery point latency. It could potentially be further behind depending on the circumstances of the mirror. If mirroring is idle and you failover, the data will only be as far behind as the recovery point latency time.

Mirror start time

The UTC time when mirroring started

Mirror end time

The UTC time when mirroring ended

Total time for last mirror

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

Validating a full server migration job

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

1. From the **Jobs** page, highlight the job and click **View Job Details** in the toolbar.
2. In the **Tasks** area on the right on the **View Job Details** page, click **Validate job properties**.
3. Carbonite Migrate 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 Carbonite Migrate correct the problem for you. For those errors that Carbonite Migrate cannot correct automatically, you will need to modify the source or target to correct the error, or you can select a different target. You must revalidate the selected servers, by clicking **Recheck**, until the validation check passes without errors.

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

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

Editing a full server migration job

Use these instructions to edit a full server migration job.

1. From the **Jobs** page, highlight the job and click **Edit Job Properties** in the toolbar. (You will not be able to edit a job if you have removed the source of that job from your Carbonite Replication Console session or if you only have Carbonite Migrate monitor security access.)
2. You will see the same options for your full server migration job as when you created the job, but you will not be able to edit all of them. If desired, edit those options that are configurable for an existing job. See *Creating a full server migration job* on page 112 for details on each job option.



Changing some options may require Carbonite Migrate to automatically disconnect, reconnect, and remirror the job.

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

-
3. If you want to modify the workload items or replication rules for the job, click **Edit workload or replication rules**. Modify the **Workload item** you are protecting, if desired. Additionally, you can modify the specific **Replication Rules** for your job.

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

You can also enter wildcard rules, however you should do so carefully. Rules are applied to files that are closest in the directory tree to them. If you have rules that include multiple folders, an exclusion rule with a wild card will need to be added for each folder that it needs applied to. For example, if you want to exclude all .log files from D:\ and your rules include D:\, D:\Dir1 , and D:\Dir2, you would need to add the exclusion rule for the root and each subfolder rule. So you will need to add exclude rules for D:*.log , D:\Dir1*.log, and D:\Dir2*.log.

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

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



If you remove data from your workload and that data has already been sent to the target, you will need to manually remove that data from the target. Because the data you removed is no longer included in the replication rules, Carbonite Migrate orphan file detection cannot remove the data for you. Therefore, you have to remove it manually.

-
4. Click **Next** to continue.
 5. Carbonite Migrate 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 Carbonite Migrate correct the problem for you. For those errors that Carbonite Migrate 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 *Carbonite Availability and Carbonite Migrate Reference Guide* for details on the various Carbonite Migrate log files.

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

Viewing a full server migration job log

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



Time	Description
6/22/2017 2:42:52 PM	Hardware IDs as follows: Source = 'bb61e75f-f6d5-4c53-9091-29017e974f2f, Target = '2d...
6/22/2017 2:42:52 PM	Completing initialization of new job 6d2bfb9d9-6a93-4eb4-8530-ca317cc7fcf9 (ALPHA to BE...
6/22/2017 2:42:53 PM	Initialization of job 6d2bfb9d9-6a93-4eb4-8530-ca317cc7fcf9 (ALPHA to BETA) complete
6/22/2017 2:42:53 PM	Changing to StoppedState from UninitializedState in response to InitializeEvent consumed...
6/22/2017 2:42:53 PM	Exited UninitializedState
6/22/2017 2:42:53 PM	Entered InitializedState
6/22/2017 2:42:53 PM	Starting monitor dfe2ee61-fde5-4afb-9ffa-17497808320c: name = FilesAndFolders_6d2bfb...
6/22/2017 2:42:53 PM	Entered StoppedState
6/22/2017 2:42:53 PM	Stopping monitor dfe2ee61-fde5-4afb-9ffa-17497808320c: Name = FilesAndFolders_6d2bfb...
6/22/2017 2:42:53 PM	Stopping share monitoring
6/22/2017 2:42:53 PM	Changing connection health to Warning
6/22/2017 2:42:53 PM	Event log entry written: '6008'.
6/22/2017 2:42:54 PM	Scheduler added new request 5b60863f-1ef2-4981-ae0d-44c0a79ead37
6/22/2017 2:42:54 PM	Deleted replication set named FilesAndFolders_6d2bfb9d96a934eb48530ca317cc7fcf9
6/22/2017 2:42:55 PM	Successfully created connection 92a61fa2-387a-4759-9fc8-60bc55141f08 connecting FilesA...
6/22/2017 2:42:55 PM	Attaching to engine connection on 172.31.206.200:6325 with following criteria:Guid = '92a...
6/22/2017 2:42:55 PM	Waiting 00:10:00 for source endpoint of 'FilesAndFolders_6d2bfb9d96a934eb48530ca317cc...
6/22/2017 2:42:55 PM	Established source endpoint of '172.31.206.200:6320' for engine connection with replicatio...
6/22/2017 2:42:55 PM	Updating failover options
6/22/2017 2:42:58 PM	The Double-Take engine is initialized.
6/22/2017 2:42:58 PM	Double-Take is NOT licensed to monitor or assume the identity of another machine.
6/22/2017 2:42:58 PM	The Double-Take engine source module is initialized.
6/22/2017 2:42:58 PM	The Double-Take engine target module is initialized.
6/22/2017 2:43:01 PM	Updating IPAddresses (Request - 5b60863f-1ef2-4981-ae0d-44c0a79ead37, WorkflowId - ...
6/22/2017 2:43:01 PM	Starting share monitoring
6/22/2017 2:43:03 PM	Changing targetActivationCode health to Ok
6/22/2017 2:43:03 PM	Event log entry written: '6004'.
6/22/2017 2:43:03 PM	Changing to ConnectedState from StoppedState in response to StartSucceededEvent consu...
6/22/2017 2:43:03 PM	Exited StoppedState
6/22/2017 2:43:03 PM	Entered ConnectedState
6/22/2017 2:43:03 PM	Subscribing to engine connection.
6/22/2017 2:43:03 PM	Changing sourceActivationCode health to Ok
6/22/2017 2:43:03 PM	Changing connection health to Ok
6/22/2017 2:43:03 PM	Changing to SynchronizedState from ConnectedState in response to MirrorCompletedEvent...
6/22/2017 2:43:03 PM	Entered ProtectingState
6/22/2017 2:43:03 PM	Starting monitor dfe2ee61-fde5-4afb-9ffa-17497808320c: name = FilesAndFolders_6d2bfb...
6/22/2017 2:43:03 PM	Entered SynchronizedState
6/22/2017 2:43:03 PM	Persisting shares
6/22/2017 2:43:03 PM	Event log entry written: '6008'.

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.

Cutting over full server migration jobs

When the migration mirror has completed, the target may or may not reboot automatically depending on your selection for **Wait for user intervention before cutover**. If you disabled user intervention, the target will reboot automatically to complete the migration process. If you enabled user intervention, when the migration mirror is complete, the status will change to **Protecting**. Use this time to complete any necessary tasks. When you are ready to complete the migration, use the following instructions to cutover.

1. On the **Jobs** page, highlight the job that you want to cutover and click **Failover, Cutover, or Recover** in the toolbar.
2. Select the type of cutover to perform.
 - **Cutover to live data**—Select this option to initiate a full, live cutover using the current data on the target. The source may be automatically shut down if it is still running, depending on your job configuration. The target will stand in for the source by rebooting and applying the source identity, including its system state, on the target. After the reboot, the target becomes the source, and the target no longer exists.
 - **Perform test cutover**—This option is not applicable to full server migration jobs.
 - **Cutover to a snapshot**—This option is not available for migration jobs.
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 cutover begins. The advantage to this option is that all of the data that the target has received will be applied before cutover 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 cutover will begin immediately. The advantage to this option is that cutover will occur immediately. The disadvantage is that any data in the target queue will be lost.
4. When you are ready to begin cutover, click **Cutover**.



If your NICs were configured for network load balancing (NLB), you will have to reconfigure that after cutover.

Right before failover occurs, Carbonite Migrate will stop all services that are not critical to Windows. If the stop command fails (perhaps because it is a blocking driver that cannot be shutdown, as is the case with some anti-virus software) or a third-party tool restarts any of these services, Carbonite Migrate may not be able to successfully failover files locked by the services. In this case, you may have to make manual modifications to your server after failover.

Some applications and hardware devices create and use software devices within the operating system, but they have the characteristics of a hardware device. For example, NIC teaming solutions are typically implemented in the operating system, however they are still designed to emulate a single piece of network hardware. In these cases, the device will not be cut over because it appears to be a hardware device.

If your source was disconnected from the network during the cutover process, the original job that was mirroring the data will still be active on that machine. Do not bring that machine back on the network. If you do, the original job will attempt to begin mirroring data again which could result in data loss. While the machine is still disconnected from the network, stop the Double-Take service, delete the file connect.sts located in the Carbonite Migrate installation directory, restart the Double-Take service, and bring the machine on the network if desired.

Because the Windows product activation is dependent on hardware, you may need to reactivate your Windows registration after cutover. The reactivation depends on several factors including service pack level, Windows edition, and your licensing type. If a target comes online after cutover with an activation failure, use the steps below appropriate for your license type. Additionally, if you are using Windows 2012, you may only have 60 minutes to complete the reactivation process until Windows activation tampering automatically shuts down your server.

- **Retail licensing**—Retail licensing allows the activation of a single operating system installation.
 1. Open the **System** applet in Windows **Control Panel**.
 2. Under **Windows activation** at the bottom of the page, click **Change product key**.
 3. Enter your retail license key. You may need access to the Internet or to call Microsoft to complete the activation.
- **MAK volume licensing**—Multiple Activation Key (MAK) licensing allows the activation of multiple operating system installations using the same activation key.
 1. View or download the Microsoft Volume Activation Deployment Guide from the Microsoft web site.
 2. Using an administrative user account, open a command prompt and follow the instructions from the deployment guide to activate MAK clients. Multiple reboots may be necessary before you can access a command prompt. You may need access to the Internet or to call Microsoft to complete the activation.
- **KMS volume licensing**—Key Management Service (KMS) licensing allows IT professionals to complete activations on their local network without contacting Microsoft.
 1. View or download the Microsoft Volume Activation Deployment Guide from the Microsoft web site.
 2. Using an administrative user account, open a command prompt and follow the instructions from the deployment guide to convert a MAK activation client to a KMS client. Multiple reboots may be necessary before you can access a command prompt.

After a cutover is complete and your target server is online as your source, when you login, you may have to specify a Windows reboot reason. You can specify any reason and continue. Additionally, you may see a prompt indicating a reboot is required because of a device change. You can disregard this error and select to reboot later.

Because Windows 64-bit has a strict driver signing policy, if you get a stop code 0x7b after failover, you may have drivers failing to load because the driver signatures are failing the

policy. In this case, reboot the server and press F8. Choose the option to not enforce the driver signing policy. If this allows the system to boot, then the problem is being caused by the cat file signature mismatch. If your system still fails to boot, contact technical support.

Chapter 7 Full server to ESX migration

Create a full server to ESX migration job when you want to migrate an entire physical server or virtual machine to an ESX target.

- *Full server to ESX migration requirements* on page 154—Full server to ESX migration includes specific requirements for this type of migration.
- *Creating a full server to ESX migration job* on page 159—This section includes step-by-step instructions for creating a full server to ESX migration job.
- *Managing and controlling full server to ESX migration jobs* on page 181—You can view status information about your full server to ESX migration job.
- *Cutting over full server to ESX migration jobs* on page 200—Use this section when you are ready to cutover from your source to your target, which will become your new source.

Full server to ESX migration requirements

Use the requirements for full server to ESX migration.

- **Source server**—The following operating systems are supported for on the source for full server to ESX jobs.
 - Windows 2016 and Server Core 2016
 - Windows 2012 R2 and Server Core 2012 R2
 - Windows 2012 and Server Core 2012
 - Windows 2008 R2 Service Pack 1 or later and Server Core 2008 R2 Service Pack 1 or later



Windows 2016 support is for the primary operating system features available in Windows 2012. New operating system features specific to Windows 2016, such as Nano Server, Windows Containers, and so on, are not supported.

Due to VMware limitations, Windows 2016 is not supported on ESX 5.5.

DNS updates are not supported for Server Core servers.

- **ESX server**—The ESX server that will host your target can be any of the following operating systems.
 - ESXi 5.5
 - ESXi 6.0
 - ESXi 6.5



The free versions of ESX restrict functionality that Carbonite Migrate requires. Therefore, you must use one of the paid editions of ESX.

- **vCenter**—vCenter is not required, but if you are using it, then you must use version 5.5 or later. If you upgrade your version of vCenter after it has been entered into the Carbonite Replication Console, you must remove and re-add the vCenter in order for the console to recognize the upgraded version.
- **vMotion**—Host vMotion is only supported if you are using vCenter. Storage vMotion is not supported.
- **Virtual recovery appliance**—The ESX server must have an existing virtual machine, known as a virtual recovery appliance, that meets the following requirements. (When you establish protection, the virtual recovery appliance will create a new virtual server, mount disks, format disks, and so on. If cutover occurs, the new virtual machine is detached from the virtual recovery appliance and powered on. Once the new virtual machine is online, it will have the identity, data, and system state of the source. Since the virtual recovery appliance maintains its own identity, it can be reused for additional cutovers.)
 - **Operating system**—The virtual recovery appliance can be any of the operating systems listed above for the source server in the following combinations.

- **2016 or 2012 R2**—If your appliance is Windows 2016 or Windows 2012 R2, your source can be any supported Windows operating system.
 - **2012 or 2008 R2**—If your appliance is Windows 2012 or Windows 2008 R2 Service Pack 1, your source can be Windows 2012 R2 or earlier.
-



Due to VMware limitations, Windows 2016 is not supported on ESX 5.5.

- **Operating system installation location**—Because VMware boots from the first bootable volume that is discovered, the operating system must be installed to SCSI controller 0, Slot 0 on the virtual recovery appliance.
 - **Carbonite Migrate**—The virtual recovery appliance must have Carbonite Migrate installed and licensed on it.
 - **Domain controllers**—If your source is a domain controller, it will start in a non-authoritative restore mode after cutover. This means that if the source was communicating with other domain controllers before cutover, it will require one of those domain controllers to be reachable after cutover so it can request updates. If this communication is not available, the domain controller will not function after cutover. If the source is the only domain controller, this is not an issue.
 - **File system**—Carbonite Migrate supports the NTFS file system. On Windows 2016 servers, ReFS is also supported. FAT and FAT32 are not supported. For detailed information on other file system capabilities, see *Mirroring and replication capabilities* on page 7.
 - **Microsoft .NET Framework**—Microsoft .NET Framework version 4.5.1 is required.
 - **System memory**—The minimum system memory on each server is 1 GB.
 - **Disk controller**—VMware Paravirtual SCSI Controllers are not supported.
 - **Disk space for program files**—This is the amount of disk space needed for the Carbonite Migrate program files. The amount depends on your operating system version and ranges from 350-500 MB.
-



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

Make sure you have additional disk space for Carbonite Migrate queuing, logging, and so on.

- **Server name**—Carbonite Migrate includes Unicode file system support, but your server name must still be in ASCII format. Additionally, all Carbonite Migrate servers and appliances must have a unique server name.
- **Time**—The clock on your Carbonite Migrate servers must be within a few minutes of each other, relative to UTC. Large time skews (more than five minutes) will cause Carbonite Migrate errors.
- **Protocols and networking**—Your servers must meet the following protocol and networking requirements.
 - Your servers must have TCP/IP with static IP addressing.
 - IPv4 is the only supported version.

- If you are using Carbonite Migrate over a WAN and do not have DNS name resolution, you will need to add the host names to the local hosts file on each server running Carbonite Migrate.
- **Network adapters**—Your source can have no more than ten NICs enabled.
- **NAT support**—Carbonite Migrate supports IP and port forwarding in NAT environments with the following caveats.
 - Only IPv4 is supported.
 - Only standalone servers are supported.
 - Make sure you have added your server to the Carbonite Replication Console using the correct public or private IP address. The name or IP address you use to add a server to the console is dependent on where you are running the console. Specify the private IP address of any servers on the same side of the router as the console. Specify the public IP address of any servers on the other side of the router as the console.
 - DNS failover and updates will depend on your configuration
 - Only the source or target can be behind a router, not both.
 - The DNS server must be routable from the target
- **Reverse lookup zone**—If you are using a DNS reverse lookup zone, then it must be Active Directory integrated. Carbonite Migrate is unable to determine if this integration exists and therefore cannot warn you during job creation if it doesn't exist.
- **DNS updates**—You can 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. Carbonite Migrate 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.

- **Windows firewall**—If you have Windows firewall enabled on your servers, there are two requirements for the Windows firewall configuration.
 - The Carbonite Migrate installation program will automatically attempt to configure ports 6320, 6325, and 6326 for Carbonite Migrate. If you cancel this step, you will have to configure the ports manually.
 - If you are using the Carbonite Replication Console to push installations out to your Windows servers, you will have to open firewall ports for WMI (Windows Management Instrumentation), which uses RPC (Remote Procedure Call). By default, RPC will use

ports at random above 1024, and these ports must be open on your firewall. RPC ports can be configured to a specific range by specific registry changes and a reboot. See the [Microsoft Knowledge Base article 154596](#) for instructions. Additionally, you will need to open firewall ports for SMB (server message block) communications which uses ports 135-139 and port 445, and you will need to open File and Printer Sharing. As an alternative, you can disable the Windows firewall temporarily until the push installations are complete.

See *Firewalls* on page 255 for instructions on handling firewalls in your environment.

- **Windows Management Instrumentation (WMI)**—Carbonite Migrate is dependent on the WMI service. If you do not use this service in your environment, contact technical support.
- **Supported configurations**—The following table identifies the supported configurations for a full server to ESX migration job.

Configuration	Description	Supported	Not Supported
One to one active/standby	You can migrate a single source to a single target host.	X	
One to one active/active	You cannot migrate a single source to a single target host where each server acts as both a source and target actively replicating data to each other.		X
Many to one	You can migrate many source servers to one target host. Replication occurs from each source to the one target host. This will consolidate your source servers to a single host server.	X	
One to many	You cannot migrate a single source to multiple target hosts.		X
Chained	You cannot migrate a single source to a single target host, where the target host then acts a source in order to send the original source to another target.		X
Single server	You cannot migrate a single source to itself.		X
Standalone to standalone	Your source and target appliance can be in a standalone to standalone configuration.	X	
Standalone to cluster	Your source and target appliance cannot be in a standalone to cluster configuration. The host where the target appliance is running can be in a cluster.		X


Configuration	Description	Supported	Not Supported
Cluster to standalone	Your source and target appliance cannot be in a cluster to standalone configuration. If your source is a standalone virtual machine, the host it is running on can be in a cluster.		X
Cluster to cluster	Your source and target appliance cannot be in a cluster to cluster configuration. If your source is a virtual machine, the host it is running on can be in a cluster. The host where the target appliance is running can be in a cluster.		X

Creating a full server to ESX migration job

Use the following instructions to migrate an entire server to a new virtual machine on an ESX server.

1. From the **Servers** page, right-click the server you want to migrate and select **Migrate**. You can also highlight a server, click **Create a New Job** in the toolbar, then select **Migrate**.
2. Choose the type of workload that you want to migrate. Under **Server Workloads**, in the **Workload types** pane, select **Full Server to ESX Migration**. In the **Workload items** pane, select the volumes on the source that you want to migrate.

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



3. By default, Carbonite Migrate selects your entire source for migration. If desired, click the **Replication Rules** heading and expand the volumes under **Folders**. You will see that Carbonite Migrate automatically excludes particular files that cannot be used during the migration. If desired, you can exclude other files that you do not want to migrate, but be careful when excluding data. Excluded volumes, folders, and/or files may compromise the integrity of your installed applications. There are some volumes, folders, and files (identified in *italics* text) that you will be unable to exclude, because they are required for migration. For example, the boot files cannot be excluded because that is where the system state information is stored.

Volumes and folders with a green highlight are included completely. Volumes and folders highlighted in light yellow are included partially, with individual files or folders included. If there is no highlight, no part of the volume or folder is included. To modify the items selected, highlight a volume, folder, or file and click **Add Rule**. Specify if you want to **Include** or **Exclude** the item. Also, specify if you want the rule to be recursive, which indicates the rule should automatically be

applied to the subdirectories of the specified path. If you do not select **Recursive**, the rule will not be applied to subdirectories.

You can also enter wildcard rules, however you should do so carefully. Rules are applied to files that are closest in the directory tree to them. If you have rules that include multiple folders, an exclusion rule with a wild card will need to be added for each folder that it needs applied to. For example, if you want to exclude all .log files from D:\ and your rules include D:\, D:\Dir1 , and D:\Dir2, you would need to add the exclusion rule for the root and each subfolder rule. So you will need to add exclude rules for D:*.log , D:\Dir1*.log, and D:\Dir2*.log.

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



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

4. Click **Next** to continue.
5. Choose your target server. This is the virtual recovery appliance on your ESX server. See *Full server to ESX migration requirements* on page 154.

Choose Target Server

Choose the server that will store the protected data.

▲ Current Servers

Servers:

Server ▲	Version	Product
BETA	8.1.0.614.0	DoubleTake Target

Show all servers

▼ Find a New Server

Diagnostics job

- **Current Servers**—This list contains the servers currently available in your console session. Servers that are not licensed for the workflow you have selected and those not applicable to the workload type you have selected will be filtered out of the list. Select your target server from the list. If the server you are looking for is not displayed, enable **Show all servers**. The servers in red are not available for the source server or workload type you have selected. Hover your mouse over an unavailable server to see a reason why this

server is unavailable.


- **Find a New Server**—If the server you need is not in the **Current Servers** list, click the **Find a New Server** heading. From here, you can specify a server along with credentials for logging in to the server. If necessary, you can click **Browse** to select a server from a network drill-down list.



If you enter the target server's fully-qualified domain name, the Carbonite Replication 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 security group.

6. Click **Next** to continue.
7. Choose the server where your target virtual recovery appliance is located. This is also the server where your replica virtual machine will be located.



- **Current VMware Servers**—This list contains the vCenter and ESX servers currently available in your console session. Select your server from the list.
- **Find a New VMware Server**—If the server you need is not in the **Current VMware Servers** list, click the **Find a New VMware Server** heading.
 - **vCenter/ESXi Server**—Select your server from the list. If your server is not in the list, manually type it in.
 - **User name**—Specify the root user or another user that has the administrator role on the specified server.
 - **Password**—Specify the password associated with the **User name** you entered.
 - **Domain**—If you are working in a domain environment, specify the **Domain**.

If your server name does not match the security certificate or the security certificate has expired, you will be prompted if you want to install the untrusted security certificate.

8. Click **Next** to continue.



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

9. You have many options available for your server migration job. Configure those options that are applicable to your environment.

Go to each page identified below to see the options available for that section of the **Set Options** page. After you have configured your options, continue with the next step on page 179.

- *General* on page 163
- *Replica Virtual Machine Location* on page 164
- *Replica Virtual Machine Configuration* on page 165
- *Replica Virtual Machine Volumes* on page 167
- *Replica Virtual Machine Network Settings* on page 169
- *Test Failover* on page 170
- *Failover Options* on page 172
- *Mirror, Verify & Orphaned Files* on page 173
- *Network Route* on page 176
- *Compression* on page 177
- *Bandwidth* on page 178

General



The screenshot shows a window titled "General" with a small upward-pointing arrow icon to the left of the title. Below the title bar, there is a label "Job name:" followed by a text input field. The input field contains the text "alpha to beta".

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

Replica Virtual Machine Location

Replica Virtual Machine Location

Select the datastore on the target ESX server that will hold the replica virtual machine:

Volume ▲	Total Size	Provisioned Space	Free Space	Owner	
EMC5	399.75 GB	141.46 GB	57.33 GB	esx51	
EMC6	399.75 GB	207.84 GB	25.88 GB	esx51	
EMC7	399.75 GB	349.34 GB	33.55 GB	esx51	

Select one of the volumes from the list to indicate the volume on the target where you want to store the configuration files for the new virtual server when it is created. The target volume must have enough **Free Space**. You can select the location of the .vmdk files under **Replica Virtual Machine Volumes**.

Replica Virtual Machine Configuration

	Source	Replica
Sockets	1	1
Cores per socket	1	1
Memory (MB)	4096	4096

Source Network Adapter	Replica Virtual Switch
Local Area Connection	InternalNetwork

- **Display name**—Specify the name of the replica virtual machine. This will be the display name of the virtual machine on the host system.
- **Hardware configuration**—Specify how you want the replica virtual machine to be created.
 - **Sockets**—Specify how many sockets to create on the new virtual machine. The number of sockets on the source is displayed to guide you in making an appropriate selection. If you select fewer sockets than the source, your clients may be impacted by slower responses.
 - **Cores per socket**—Specify how many cores to create per socket. The number of cores per socket on the source is displayed to guide you in making an appropriate selection.
 - **Memory**—Specify the amount of memory, in MB, to create on the new virtual machine. The memory on the source is displayed to guide you in making an appropriate selection. If you select less memory than the source, your clients may be impacted by slower responses.
- **Network adapter type**—If your target appliance has VMware Tools installed, you can select the type of adapter, **E1000** or **VmxNet3**, to use on the replica virtual machine. This selection will apply to all adapters on the replica.



If your source has VMware Tools installed, but it is an older version than VMware Tools installed on your target appliance, you will have to update VMware Tools on the replica server after failover in order to get the VmxNet3 adapter to function.

- **Virtual switches**—Identify how you want to handle the network mapping after cutover. The **Source Network Adapter** column lists the NICs from the source. Map each one to a **Replica Virtual Switch**, which is a virtual network on the target. You can also choose to discard the source's NIC and IP addresses.

- **Power on replica after failover**—By default, the replica virtual machine will automatically be powered on after the cutover process is complete. If you want the replica virtual machine to remain powered off, disable this option.

Replica Virtual Machine Volumes

Volume	Disk Size	Used Space	Replica Disk Size	Replica Disk Format	Target Datastore	Virtual Disk	Pre-existing Disk Path
C:	145.9 GB	15.92 GB	145.9 GB	Flat Disk	EMC8	Create new disk	

Notes:
Changes to the disk size and disk type will not be used for pre-existing disks because the pre-existing configuration will be used.
The pre-existing disks might have the below format.

alpha_C.vmdk

- **Replica Disk Size**—For each volume you are protecting, specify the size of the replica disk on the target. Be sure and include the value in MB or GB for the disk. The value must be at least the size of the specified **Used Space** on that volume.



In some cases, the replica virtual machine may use more virtual disk space than the size of the source volume due to differences in how the virtual disk's block size is formatted. To avoid this issue, make sure the replica can accommodate not just the size of all of the files but the size on disk as well.

- **Replica Disk Format**—For each volume you are protecting, specify the format of the disk that will be created.
 - **Flat Disk**—This disk format allocates the full amount of the disk space immediately, but does not initialize the disk space to zero until it is needed.
 - **Thick**—This disk format allocates the full amount of the disk space immediately, initializing all of the allocated disk space to zero.
 - **Thin**—This disk format does not allocate the disk space until it is needed.
- **Target Datastore**—For each volume you are protecting, specify the datastore on the target where you want to store the virtual disk files for the new replica virtual machine. You can specify the location of the virtual machine configuration files under **Replica Virtual Machine Location**. If you are going to reuse an existing disk, select the volume where the disk is located.
- **Virtual Disk**—Specify if you want Carbonite Migrate to create a new disk for your replica virtual machine or if you want to use an existing disk.

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

Each pre-existing disk must be located on the target datastore specified. If you have copied the .vmdk file to this location manually, be sure you have also copied the associated -flat.vmdk file too. If you have used vCenter to copy the virtual machine, the associated file will automatically be copied. There are no restrictions on the file name of the .vmdk, but the

associated -flat.vmdk file must have the same base name and the reference to that flat file in the .vmdk must be correct. Carbonite Migrate will move, not copy, the virtual disk files to the appropriate folders created by the replica, so make sure the selected target datastore is where you want the replica virtual disk to be located.

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

- a. Create a job in a LAN environment, letting Carbonite Migrate create the virtual disk for you.
 - b. Complete the mirror process locally.
 - c. Delete the job and when prompted, do not delete the replica.
 - d. Remove the replica virtual machine from the ESX inventory, which will delete the virtual machine configuration but will keep the associated .vmdk files.
 - e. Located the .vmdk files on your datastore that were left behind when you removed the virtual machine from the inventory. Rename the folder where the .vmdk files are located.
 - f. Shut down and move the ESX target server to your remote site.
 - g. After the ESX target server is back online at the remote site, move the .vmdk files to a temporary location.
 - h. Create a new protection job for the same source server and select to **Use existing disk**, specifying the temporary location of your .vmdk files. Carbonite Migrate will reuse the existing .vmdk files (automatically moving the files to the correct location) and perform a difference mirror over the WAN to bring the virtual machine up-to-date.
- **Pre-existing Disk Path**—This is the location of your existing virtual disks on the selected **Target Volume** that you want to reuse.

Replica Virtual Machine Network Settings

Use advanced settings for replica virtual machine network configuration.

Network adapters:
Local Area Connection (112.42.74.29)

Source IP addresses:

IP Address	Subnet Mask
112.42.74.29	255.255.0.0

Replica IP addresses:

IP Address	Subnet Mask
112.52.74.29	255.255.0.0

Source Default Gateways:

112.42.48.9

Replica Default Gateways:

112.52.48.9

Source DNS Server addresses:

112.42.48.20

Replica DNS Server addresses:

112.52.48.20

- **Use advanced settings for replica virtual machine network configuration**—Select this option to enable the replica virtual machine network setting configuration. This setting is primarily used for WAN support.
- **Network adapters**—Select a network adapter from the source and specify the **Replica IP addresses**, **Replica Default Gateways**, and **Replica DNS Server addresses** to be used after cutover. 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 cutover 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 cutover.

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.

Test Failover

These options allow you to perform a test failover. Keep in mind the following for using test failover.

- The source, target, and protection job will remain online and uninterrupted during the test.
- The test will be performed using the test failover settings configured during job creation.
- The test will use the current data on the target.
- The test failover will take a snapshot of the current data on the target and create a new set of virtual disks. The data from the snapshot will be mirrored to the new set of disks using the same mirroring options as the protection job.
- Once the mirror is complete the replica virtual machine is automatically brought online using the new set of disks.
- The replica virtual machine will use the network settings specified in the test failover settings of the protection job.
- When you are finished with your test, undo it.
- When you undo a test failover, the new set of disks will be maintained or deleted as specified in the test failover settings of the protection job.
- At any time during a test failover, you can undo the test, perform a live failover, or failover to a snapshot. (Performing a live failover or failing over to a snapshot will automatically undo any test in progress.)

Source Network Adapter	Replica Virtual Switch
Local Area Connection	VM Network 5

Volume	Replica Disk Format	Target Datastore
C:	Flat disk	RNDESX65-DS1
E:	Thin	RNDESX65-DS1

- **Do not connect replica network adapters on test failover**—Select this option if you do not want the replica virtual machine used for the test to be connected to the network.
- **Connect and map replica network adapters on test failover**—Select this option if you want the replica virtual machine used for the test to be connected to the network. You will need to map each **Source Network Adapter** to a **Target Virtual Switch** for the test. You can also choose to discard the source's NIC and IP addresses.
- **Replica Disk Format**—For each volume you are protecting, specify the format of the disk that will be created on the test server.

- **Flat Disk**—This disk format allocates the full amount of the disk space immediately, but does not initialize the disk space to zero until it is needed.
 - **Thick**—This disk format allocates the full amount of the disk space immediately, initializing all of the allocated disk space to zero.
 - **Thin**—This disk format does not allocate the disk space until it is needed.
 - **Target Datastore**—For each volume you are protecting, specify the datastore on the target where you want to store the virtual disk files for the test server.
 - **Delete test failover virtual disks**—Select this option if you want to delete the new virtual disks created during the test failover process. If you disable this option, the new disks will not be deleted when you perform undo failover.
-



Be careful if you choose to connect the network adapters for a test cutover. Depending on your network adapter mappings, users may be able to access the target. Also, since the source is still online, there is a chance users may split between accessing the source or target.

Failover Options



Failover Options

- Wait for user to initiate failover
- Shutdown source server

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

Mirror, Verify & Orphaned Files

Mirror, Verify & Orphaned Files

Mirror Options

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

Compare file attributes. Send the attributes and bytes that differ.

Verification Options

Enable scheduled verification

Verify on this interval: 1 Days

Begin immediately

Begin at this time: 3/23/2017 10:53:32 AM

Report and comparison options

Report only

Report and mirror files

Compare file attributes and data

General Options

Calculate size of protected data upon connection

Delete orphaned files

- **Mirror Options**—Choose a comparison method and whether to mirror the entire file or only the bytes that differ in each file.
 - **Do not compare files. Send the entire file.**—Carbonite Migrate 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 option requires no time for comparison, but it can be slower because it sends the entire file. However, it is useful for configurations that have large data sets with millions of small files that are frequently changing and it is more efficient to send the entire file. You may also need to use this option if configuration management policies require sending the entire file.
 - **Compare file attributes. Send the attributes and bytes that differ.**—Carbonite Migrate will compare file attributes and will mirror only the attributes and bytes that are different. This option is the fastest comparison method and fastest mirror option. Files that have not changed can be easily skipped. Also files that are open and require a checksum mirror can be compared.
 - **Compare file attributes and data. Send the attributes and bytes that differ.**—Carbonite Migrate will compare file attributes and the file data and will mirror only the attributes and bytes that are different. This comparison method is not as fast because every file is compared, regardless of whether the file has changed or is open. However, sending only the attributes and bytes that differ is the fastest mirror 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 Carbonite Migrate



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, Carbonite Migrate 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 Carbonite Migrate 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.

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

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

If you want to move orphaned files rather than delete them, you can configure this option along with the move deleted files feature to move your orphaned files to the specified deleted files directory. See *Target server properties* on page 47 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 Carbonite Migrate 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 Carbonite Migrate log, and during difference mirrors, verifications, and restorations, all orphaned file processing messages are logged to the Carbonite Migrate log. The orphaned file log is located in the **Logging folder** specified for the source. See *Log file properties* on page 52 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




Send data to the target server using this route:
10.10.10.30

By default, Carbonite Migrate 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 Carbonite Migrate traffic. For example, you can separate regular network traffic and Carbonite Migrate traffic on a machine with multiple IP addresses. You can also select or manually enter a public IP address (which is the public IP address of the server's router) if you are using a NAT environment. The **Management Service port** and **Replication Service port** can be disregarded. It is used for other job types.



The IP address used on the source will be determined through the Windows route table.

Compression



To help reduce the amount of bandwidth needed to transmit Carbonite Migrate 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 Carbonite Migrate data.



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

Bandwidth



Bandwidth


Do not limit bandwidth

Use a fixed limit (bytes per second):

Preset bandwidth: Bandwidth (bytes per second):

Use scheduled limits:

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

 Bandwidth is not limited during unscheduled times.

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



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

- **Do not limit bandwidth**—Carbonite Migrate will transmit data using 100% bandwidth availability.
- **Use a fixed limit**—Carbonite Migrate 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**—Carbonite Migrate 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**.

10. Click **Next** to continue.
11. Carbonite Migrate 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 Carbonite Migrate correct the problem for you. For those errors that Carbonite Migrate 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 *Carbonite Availability and Carbonite Migrate Reference Guide* for details on the various Carbonite Migrate log files.

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



Jobs in a NAT environment may take longer to start.

Once a job is created, do not change the name of underlying hardware components used in the job. For example, volume and datastore names or network adapter and virtual switch names. Any component used by name in your job must continue to use that name throughout the lifetime of job. If you must change a name, you will need to delete the job and re-create it using the new component name.

Managing and controlling full server to ESX migration jobs

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

The jobs displayed in the right pane depend on the server group folder selected in the left pane. Every job for each server in your console session is displayed when the **Jobs on All Servers** group is selected. If you have created and populated server groups (see *Managing servers* on page 18), then only the jobs associated with the server or target servers in that server group will be displayed in the right pane.

- *Overview job information displayed in the top right pane* on page 181
- *Detailed job information displayed in the bottom right pane* on page 184
- *Job controls* on page 186


Overview job information displayed in the top right pane


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


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


Column 1 (Blank)

The first blank column indicates the state of the job.

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

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

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

 The job is in an unknown state.

Job

The name of the job

Source Server

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

Target Server

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

Job Type

Each job type has a unique job type name. This job is a Full Server to ESX Migration job. For a complete list of all job type names, press F1 to view the Carbonite Replication Console online help.

Activity

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

Mirror Status

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

Replication Status

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

Transmit Mode

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

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

Operating System

The job type operating system

Detailed job information displayed in the bottom right pane

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

Name

The name of the job

Target data state

- **OK**—The data on the target is in a good state.
- **Mirroring**—The target is in the middle of a mirror process. The data will not be in a good state until the mirror is complete.
- **Mirror Required**—The data on the target is not in a good state because a remirror is required. This may be caused by an incomplete or stopped mirror or an operation may have been dropped on the target.
- **Busy**—The source is low on memory causing a delay in getting the state of the data on the target.
- **Not Loaded**—Carbonite Migrate 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 38.

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 38.

Replication queue

The total number of replication bytes in the source queue

Disk queue

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

Recovery point latency

The length of time replication is behind on the target compared to the source. This is the time period of replication data that would be lost if a failure were to occur at the current time. This value represents replication data only and does not include mirroring data. If you are mirroring and failover, the data on the target will be at least as far behind as the recovery point latency. It could potentially be further behind depending on the

circumstances of the mirror. If mirroring is idle and you failover, the data will only be as far behind as the recovery point latency time.

Bytes sent

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

Bytes sent (compressed)

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

Connected since

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

Recent activity

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

Additional information

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

Job controls

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

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

View Job Details

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

Edit Job Properties

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

Delete

Stops (if running) and deletes the selected jobs.

Provide Credentials

Changes the login credentials that the job (which is on the target machine) uses to authenticate to the servers in the job. This button opens the Provide Credentials dialog box where you can specify the new account information and which servers you want to update. See *Providing server credentials* on page 31. You will remain on the **Jobs** page after updating the server credentials. If your servers use the same credentials, make sure you also update the credentials on the **Servers** page so that the Carbonite Replication Console can authenticate to the servers in the console session. See *Managing servers* on page 18.

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 Carbonite Migrate queue was not exhausted during

the time the job was paused. If the Carbonite Migrate queue was exhausted during the time the job was paused, the job will restart mirroring and replication.

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

Pause

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

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

Stop

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

Take Snapshot

Snapshots are not applicable to migration jobs.

Manage Snapshots

Snapshots are not applicable to migration jobs.

Failover or Cutover

Starts the cutover process. See *Cutting over full server to ESX migration jobs* on page 200 for the process and details of cutting over a full server to ESX migration job.

Failback

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

Restore

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

Reverse

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

Undo Failover or Cutover

Cancels a test cutover by undoing it. This resets the servers and the job back to their original state. See *Cutting over full server to ESX migration jobs* on page 200 for the process and details of undoing a cut over full server to ESX migration 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, Carbonite Migrate 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, Carbonite Migrate 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.**—Carbonite Migrate 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 option requires no time for comparison, but it can be slower because it sends the entire file. However, it is useful for configurations that have large data sets with millions of small files that are frequently changing and it is more efficient to send the entire file. You may also need to use this option if configuration management policies require sending the entire file.
 - **Compare file attributes. Send the attributes and bytes that differ.**—Carbonite Migrate will compare file attributes and will mirror only the attributes and bytes that are different. This option is the fastest comparison method and fastest mirror option. Files that have not changed can be easily skipped. Also files that are open and require a checksum mirror can be compared.
 - **Compare file attributes and data. Send the attributes and bytes that differ.**—Carbonite Migrate will compare file attributes and the file data and will mirror only the attributes and bytes that are different. This comparison method is not as fast because every file is compared, regardless of whether the file has changed or is open. However, sending only the attributes and bytes that differ is the fastest mirror option.
- **Calculate size of protected data before mirroring**—Specify if you want

Carbonite Migrate 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**—Carbonite Migrate will transmit data using 100% bandwidth availability.
 - **Fixed bandwidth limit**—Carbonite Migrate 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 Carbonite Migrate 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 Carbonite Migrate processing, not the entire server.

While the target is paused, the Carbonite Migrate 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, Carbonite Migrate will automatically disconnect the connections and attempt to reconnect them.

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

Filter

Select a filter option from the drop-down list to only display certain jobs. You can display **Healthy jobs**, **Jobs with warnings**, or **Jobs with errors**. To clear the filter, select **All jobs**. If you have created and populated server groups, then the filter will only apply to the jobs associated with the server or target servers in that server group. See *Managing servers* on page 18.

Search

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

Overflow Chevron

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

Viewing full server to ESX migration job details

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

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





Job name

The name of the job

Job type

Each job type has a unique job type name. This job is a Full Server to ESX Migration job. For a complete list of all job type names, press F1 to view the Carbonite Replication Console online help.

Health

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

Activity

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

Connection ID

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

Transmit mode

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

Target data state

- **OK**—The data on the target is in a good state.
- **Mirroring**—The target is in the middle of a mirror process. The data will not be in a good state until the mirror is complete.
- **Mirror Required**—The data on the target is not in a good state because a remirror is required. This may be caused by an incomplete or stopped mirror or an operation may have been dropped on the target.
- **Busy**—The source is low on memory causing a delay in getting the state of the data on the target.
- **Not Loaded**—Carbonite Migrate 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 Carbonite Migrate transmissions.

Compression

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

Encryption

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

Bandwidth limit

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

Connected since

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

Additional information

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

Mirror status

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

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

Mirror percent complete

The percentage of the mirror that has been completed

Mirror remaining

The total number of mirror bytes that are remaining to be sent from the source to the target. This value may be zero if you have enabled **Mirror only changed files when source reboots** on the *Server setup properties* on page 38.

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 38.

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 Carbonite Migrate driver. Check the Event Viewer for driver related issues.
- **Unknown**—The console cannot determine the status.

Replication queue

The total number of replication bytes in the source queue

Disk queue

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

Bytes sent

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

Bytes sent compressed

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

Recovery point latency

The length of time replication is behind on the target compared to the source. This is the time period of replication data that would be lost if a failure were to occur at the current time. This value represents replication data only and does not include mirroring data. If you are mirroring and failover, the data on the target will be at least as far behind as the recovery point latency. It could potentially be further behind depending on the circumstances of the mirror. If mirroring is idle and you failover, the data will only be as far behind as the recovery point latency time.

Mirror start time

The UTC time when mirroring started

Mirror end time

The UTC time when mirroring ended

Total time for last mirror

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

Validating a full server to ESX migration job

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

1. From the **Jobs** page, highlight the job and click **View Job Details** in the toolbar.
2. In the **Tasks** area on the right on the **View Job Details** page, click **Validate job properties**.
3. Carbonite Migrate 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 Carbonite Migrate correct the problem for you. For those errors that Carbonite Migrate cannot correct automatically, you will need to modify the source or target to correct the error, or you can select a different target. You must revalidate the selected servers, by clicking **Recheck**, until the validation check passes without errors.

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

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

Editing a full server to ESX migration job

Use these instructions to edit a full server to ESX migration job.

1. From the **Jobs** page, highlight the job and click **Edit Job Properties** in the toolbar. (You will not be able to edit a job if you have removed the source of that job from your Carbonite Replication Console session or if you only have Carbonite Migrate monitor security access.)
2. You will see the same options available for your full server to ESX migration job as when you created the job, but you will not be able to edit all of them. If desired, edit those options that are configurable for an existing job. See *Creating a full server to ESX migration job* on page 159 for details on each job option.



Changing some options may require Carbonite Migrate to automatically disconnect, reconnect, and remirror the job.

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

3. If you want to modify the workload items or replication rules for the job, click **Edit workload or replication rules**. Modify the **Workload item** you are protecting, if desired. Additionally, you can modify the specific **Replication Rules** for your job.

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

You can also enter wildcard rules, however you should do so carefully. Rules are applied to files that are closest in the directory tree to them. If you have rules that include multiple folders, an exclusion rule with a wild card will need to be added for each folder that it needs applied to. For example, if you want to exclude all .log files from D:\ and your rules include D:\, D:\Dir1 , and D:\Dir2, you would need to add the exclusion rule for the root and each subfolder rule. So you will need to add exclude rules for D:*.log , D:\Dir1*.log, and D:\Dir2*.log.

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

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



If you remove data from your workload and that data has already been sent to the target, you will need to manually remove that data from the target. Because the data you

removed is no longer included in the replication rules, Carbonite Migrate orphan file detection cannot remove the data for you. Therefore, you have to remove it manually.

4. Click **Next** to continue.
5. Carbonite Migrate 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 Carbonite Migrate correct the problem for you. For those errors that Carbonite Migrate 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 *Carbonite Availability and Carbonite Migrate Reference Guide* for details on the various Carbonite Migrate log files.

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

Viewing a full server to ESX migration job log

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



Time	Description
6/22/2017 2:42:52 PM	Hardware IDs as follows: Source = 'bb61e75f-f6d5-4c53-9091-29017e974f2f', Target = '2d...
6/22/2017 2:42:52 PM	Completing initialization of new job 6d2bfb9d9-6a93-4eb4-8530-ca317cc7fcf9 (ALPHA to BE...
6/22/2017 2:42:53 PM	Initialization of job 6d2bfb9d9-6a93-4eb4-8530-ca317cc7fcf9 (ALPHA to BETA) complete
6/22/2017 2:42:53 PM	Changing to StoppedState from UninitializedState in response to InitializeEvent consumed...
6/22/2017 2:42:53 PM	Exited UninitializedState
6/22/2017 2:42:53 PM	Entered InitializedState
6/22/2017 2:42:53 PM	Starting monitor dfe2ee61-fde5-4afb-9ffa-17497808320c: name = FilesAndFolders_6d2bfb...
6/22/2017 2:42:53 PM	Entered StoppedState
6/22/2017 2:42:53 PM	Stopping monitor dfe2ee61-fde5-4afb-9ffa-17497808320c: Name = FilesAndFolders_6d2bfb...
6/22/2017 2:42:53 PM	Stopping share monitoring
6/22/2017 2:42:53 PM	Changing connection health to Warning
6/22/2017 2:42:53 PM	Event log entry written: '6008'.
6/22/2017 2:42:54 PM	Scheduler added new request 5b60863f-1ef2-4981-ae0d-44c0a79ead37
6/22/2017 2:42:54 PM	Deleted replication set named FilesAndFolders_6d2bfb9d96a934eb48530ca317cc7fcf9
6/22/2017 2:42:55 PM	Successfully created connection 92a61fa2-387a-4759-9fc8-60bc55141f08 connecting FilesA...
6/22/2017 2:42:55 PM	Attaching to engine connection on 172.31.206.200:6325 with following criteria:Guid = '92a...
6/22/2017 2:42:55 PM	Waiting 00:10:00 for source endpoint of 'FilesAndFolders_6d2bfb9d96a934eb48530ca317cc...
6/22/2017 2:42:55 PM	Established source endpoint of '172.31.206.200:6320' for engine connection with replicatio...
6/22/2017 2:42:55 PM	Updating failover options
6/22/2017 2:42:58 PM	The Double-Take engine is initialized.
6/22/2017 2:42:58 PM	Double-Take is NOT licensed to monitor or assume the identity of another machine.
6/22/2017 2:42:58 PM	The Double-Take engine source module is initialized.
6/22/2017 2:42:58 PM	The Double-Take engine target module is initialized.
6/22/2017 2:43:01 PM	Updating IPAddresses (Request - 5b60863f-1ef2-4981-ae0d-44c0a79ead37, WorkflowId - ...
6/22/2017 2:43:01 PM	Starting share monitoring
6/22/2017 2:43:03 PM	Changing targetActivationCode health to Ok
6/22/2017 2:43:03 PM	Event log entry written: '6004'.
6/22/2017 2:43:03 PM	Changing to ConnectedState from StoppedState in response to StartSucceededEvent consu...
6/22/2017 2:43:03 PM	Exited StoppedState
6/22/2017 2:43:03 PM	Entered ConnectedState
6/22/2017 2:43:03 PM	Subscribing to engine connection.
6/22/2017 2:43:03 PM	Changing sourceActivationCode health to Ok
6/22/2017 2:43:03 PM	Changing connection health to Ok
6/22/2017 2:43:03 PM	Changing to SynchronizedState from ConnectedState in response to MirrorCompletedEvent...
6/22/2017 2:43:03 PM	Entered ProtectingState
6/22/2017 2:43:03 PM	Starting monitor dfe2ee61-fde5-4afb-9ffa-17497808320c: name = FilesAndFolders_6d2bfb...
6/22/2017 2:43:03 PM	Entered SynchronizedState
6/22/2017 2:43:03 PM	Persisting shares
6/22/2017 2:43:03 PM	Event log entry written: '6008'.

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.

Cutting over full server to ESX migration jobs

When the migration mirror has completed, the target may or may not reboot automatically depending on your selection for **Wait for user intervention before cutover**. If you disabled user intervention, the target will reboot automatically to complete the migration process. If you enabled user intervention, when the migration mirror is complete, the status will change to **Protecting**. Use this time to complete any necessary tasks. When you are ready to complete the migration, use the following instructions to cutover.

1. On the **Jobs** page, highlight the job that you want to cutover and click **Failover, Cutover, or Recover** in the toolbar.
2. Select the type of cutover to perform.
 - **Cutover to live data**—Select this option to initiate a full, live cutover using the current data on the target. The source may be automatically shut down if it is still running, depending on your job configuration. The protection job is stopped and the replica virtual machine is started on the target with full network connectivity.
 - **Perform test cutover**—Select this option to perform a test cutover.
 - The source, target, and protection job will remain online and uninterrupted during the test.
 - The test will be performed using the test failover settings configured during job creation.
 - The test will use the current data on the target.
 - The test failover will take a snapshot of the current data on the target and create a new set of virtual disks. The data from the snapshot will be mirrored to the new set of disks using the same mirroring options as the protection job.
 - Once the mirror is complete the replica virtual machine is automatically brought online using the new set of disks.
 - The replica virtual machine will use the network settings specified in the test failover settings of the protection job.
 - When you are finished with your test, undo it.
 - When you undo a test failover, the new set of disks will be maintained or deleted as specified in the test failover settings of the protection job.
 - At any time during a test failover, you can undo the test, perform a live failover, or failover to a snapshot. (Performing a live failover or failing over to a snapshot will automatically undo any test in progress.)
 - **Cutover to a snapshot**—This option is not available for migration jobs.
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 cutover begins. The advantage to this option is that all of the data that the target has received will be applied before cutover 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 cutover will begin immediately. The advantage to this option is that cutover will occur immediately. The disadvantage is that any data in the target queue will be lost.

4. When you are ready to begin cutover, click **Cutover**.



Once cutover has started, do not reboot the target appliance. If the cutover process is interrupted, it may fail.



If your source was disconnected from the network during the cutover process, the original job that was mirroring the data will still be active on that machine. Do not bring that machine back on the network. If you do, the original job will attempt to begin mirroring data again which could result in data loss. While the machine is still disconnected from the network, stop the Double-Take service, delete the file connect.sts located in the Carbonite Migrate installation directory, restart the Double-Take service, and bring the machine on the network if desired.

Because the Windows product activation is dependent on hardware, you may need to reactivate your Windows registration after cutover. The reactivation depends on several factors including service pack level, Windows edition, and your licensing type. If a target comes online after cutover with an activation failure, use the steps below appropriate for your license type. Additionally, if you are using Windows 2012, you may only have 60 minutes to complete the reactivation process until Windows activation tampering automatically shuts down your server.

- **Retail licensing**—Retail licensing allows the activation of a single operating system installation.
 1. Open the **System** applet in Windows **Control Panel**.
 2. Under **Windows activation** at the bottom of the page, click **Change product key**.
 3. Enter your retail license key. You may need access to the Internet or to call Microsoft to complete the activation.
- **MAK volume licensing**—Multiple Activation Key (MAK) licensing allows the activation of multiple operating system installations using the same activation key.
 1. View or download the Microsoft Volume Activation Deployment Guide from the Microsoft web site.
 2. Using an administrative user account, open a command prompt and follow the instructions from the deployment guide to activate MAK clients. Multiple reboots may be necessary before you can access a command prompt. You may need access to the Internet or to call Microsoft to complete the activation.
- **KMS volume licensing**—Key Management Service (KMS) licensing allows IT professionals to complete activations on their local network without contacting Microsoft.
 1. View or download the Microsoft Volume Activation Deployment Guide from the Microsoft web site.
 2. Using an administrative user account, open a command prompt and follow the instructions from the deployment guide to convert a MAK activation client to a KMS client. Multiple reboots may be necessary before you can access a command prompt.

If your job was using vCenter, you may have problems with cutover if vCenter is down or if it is unreachable. Contact technical support for details on how to complete cutover in this situation.

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

After failover is complete, an additional NIC may be seen when looking at hidden devices or a system scan. These NICs are not loaded so they are not consuming any resources and are not active. You can safely disregard these hidden NICs.

5. If you performed a test cutover, you can undo it by selecting **Undo Failover or Cutover** in the toolbar. The replica virtual machine on the target will be shut down and, if configured, the virtual disks used for the test cutover will be deleted.

Chapter 8 Full server to Hyper-V migration

Create a full server to Hyper-V migration job when you want to migrate an entire physical server or virtual machine to a Hyper-V target.

- *Full server to Hyper-V migration requirements* on page 204—Full server to Hyper-V migration includes specific requirements for this type of migration.
- *Creating a full server to Hyper-V migration job* on page 209—This section includes step-by-step instructions for creating a full server to Hyper-V migration job.
- *Managing and controlling full server to Hyper-V migration jobs* on page 231—You can view status information about your full server to Hyper-V migration job.
- *Cutting over full server to Hyper-V migration jobs* on page 250—Use this section when you are ready to cutover from your source to your target, which will become your new source.

Full server to Hyper-V migration requirements

Use the following requirements for full server to Hyper-V migration.

- **Source server**—The following operating systems are supported for on the source for full server to Hyper-V jobs.
 - Windows 2016 and Server Core 2016
 - Windows 2012 R2 and Server Core 2012 R2
 - Windows 2012 and Server Core 2012
 - Windows 2008 R2 Service Pack 1 or later and Server Core 2008 R2 Service Pack 1 or later



Windows 2016 support is for the primary operating system features available in Windows 2012. New operating system features specific to Windows 2016, such as Nano Server, Windows Containers, and so on, are not supported.

Due to VMware limitations, Windows 2016 is not supported on ESX 5.5.

DNS updates are not supported for Server Core servers.

-
- **Target host**—Your physical Hyper-V host machine can be any operating system from the source requirements above that has the Hyper-V role enabled. In addition, you can use Hyper-V Server 2008 R2, or any of the Server Core operating systems with the Hyper-V role enabled. (Hyper-V Server 2008 and Server Core 2008 are not supported.)



The target host can only be used in the following combinations.

- **2016 or 2012 R2**—If your host is Windows 2016 or Windows 2012 R2, your source can be any supported Windows operating system.
- **2012 or 2008 R2**—If your host is Windows 2012 or Windows 2008 R2 Service Pack 1, you source can be Windows 2012 R2 or earlier.

-
- **File system**—Carbonite Migrate supports the NTFS file system. On Windows 2016 servers, ReFS is also supported. FAT and FAT32 are not supported. For detailed information on other file system capabilities, see *Mirroring and replication capabilities* on page 7.
 - **Microsoft .NET Framework**—Microsoft .NET Framework version 4.5.1 is required.
 - **System memory**—The minimum system memory on each server is 1 GB.
 - **Disk types**—Source virtual machines can use raw, pass-through, or differencing disks, however, they will be virtual hard disks on the replica on the target.
 - **Disk space for program files**—This is the amount of disk space needed for the Carbonite Migrate program files. The amount depends on your operating system version and ranges from 350-500 MB.



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

Make sure you have additional disk space for Carbonite Migrate queuing, logging, and so on.

- **Server name**—Carbonite Migrate includes Unicode file system support, but your server name must still be in ASCII format. Additionally, all Carbonite Migrate servers and appliances must have a unique server name.
- **Time**—The clock on your Carbonite Migrate servers must be within a few minutes of each other, relative to UTC. Large time skews (more than five minutes) will cause Carbonite Migrate errors.
- **Protocols and networking**—Your servers must meet the following protocol and networking requirements.
 - Your servers must have TCP/IP with static IP addressing.
 - IPv4 is the only supported version.
 - If you are using Carbonite Migrate over a WAN and do not have DNS name resolution, you will need to add the host names to the local hosts file on each server running Carbonite Migrate.
- **Network adapters**—Your source can have no more than 12 NICs enabled (eight synthetic and four legacy).
- **NAT support**—Carbonite Migrate supports IP and port forwarding in NAT environments with the following caveats.
 - Only IPv4 is supported.
 - Only standalone servers are supported.
 - Make sure you have added your server to the Carbonite Replication Console using the correct public or private IP address. The name or IP address you use to add a server to the console is dependent on where you are running the console. Specify the private IP address of any servers on the same side of the router as the console. Specify the public IP address of any servers on the other side of the router as the console.
 - DNS failover and updates will depend on your configuration
 - Only the source or target can be behind a router, not both.
 - The DNS server must be routable from the target
- **Reverse lookup zone**—If you are using a DNS reverse lookup zone, then it must be Active Directory integrated. Carbonite Migrate is unable to determine if this integration exists and therefore cannot warn you during job creation if it doesn't exist.
- **DNS updates**—You can 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. Carbonite Migrate 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.

- **Windows firewall**—If you have Windows firewall enabled on your servers, there are two requirements for the Windows firewall configuration.
 - The Carbonite Migrate installation program will automatically attempt to configure ports 6320, 6325, and 6326 for Carbonite Migrate. If you cancel this step, you will have to configure the ports manually.
 - If you are using the Carbonite Replication Console to push installations out to your Windows servers, you will have to open firewall ports for WMI (Windows Management Instrumentation), which uses RPC (Remote Procedure Call). By default, RPC will use ports at random above 1024, and these ports must be open on your firewall. RPC ports can be configured to a specific range by specific registry changes and a reboot. See the [Microsoft Knowledge Base article 154596](#) for instructions. Additionally, you will need to open firewall ports for SMB (server message block) communications which uses ports 135-139 and port 445, and you will need to open File and Printer Sharing. As an alternative, you can disable the Windows firewall temporarily until the push installations are complete.

See *Firewalls* on page 255 for instructions on handling firewalls in your environment.

- **Windows Management Instrumentation (WMI)**—Carbonite Migrate is dependent on the WMI service. If you do not use this service in your environment, contact technical support.
- **Supported configurations**—The following table identifies the supported configurations for a full server to Hyper-V migration job.

Configuration	Description	Supported	Not Supported
One to one active/standby	You can migrate a single source to a single target host.	X	
One to one active/active	You cannot migrate a single source to a single target host where each server acts as both a source and target actively replicating data to each other.		X
Many to one	You can migrate many source servers to one target host. Replication occurs from each source to the one target host. This will consolidate your source servers to a single host server.	X	
One to many	You cannot migrate a single source to multiple target hosts.		X

Configuration	Description	Supported	Not Supported
Chained	You cannot migrate a single source to a single target host, where the target host then acts a source in order to send the original source to another target.		X
Single server	You cannot migrate a single source to itself.		X
Standalone to standalone	Your source and target host can be in a standalone to standalone configuration.	X	
Standalone to cluster	<p>Your source and target host can be in a standalone to cluster configuration. The target cluster will be cluster-aware and will roll and failover to cluster nodes using standard Microsoft clustering. If you are using a target cluster, make sure your cluster meets the following requirements.</p> <ul style="list-style-type: none"> • Carbonite Migrate disk queue— Ensure that the disk queue is on a locally attached, non-cluster volume. • Licensing—Each node in the cluster must have a valid Carbonite Availability license key. • Resource registration—In some cases, the Carbonite Migrate cluster resources may not be registered automatically when Carbonite Migrate 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. 	X	
Cluster to standalone	Your source and target cannot be in a cluster to standalone configuration. However, if your source is a standalone virtual machine, the host it is running on can be in a cluster.		X
Cluster to cluster	Your source and target cannot be in a cluster to cluster configuration. However, if your source is a standalone virtual machine, the host it is running on can be in a cluster. The target replica will be cluster-aware and will roll and failover to cluster nodes using standard Microsoft clustering.		X

If you are using a Cluster Shared Volume (CSV), you cannot migrate any data, including a virtual machine, residing on the CSV. If you want to migrate a CSV virtual machine, you must run Carbonite Migrate from within the guest operating system of the virtual machine and create the job within the guest.


Data can be written to a target CSV.

Creating a full server to Hyper-V migration job

Use the following instructions to create a full server to Hyper-V migration job.

1. From the **Servers** page, right-click the server you want to migrate and select **Migrate**. You can also highlight a server, click **Create a New Job** in the toolbar, then select **Migrate**.
2. Choose the type of workload that you want to migrate. Under **Server Workloads**, in the **Workload types** pane, select **Full Server to Hyper-V Migration**. In the **Workload items** pane, select the volumes on the source that you want to migrate.

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



3. By default, Carbonite Migrate selects your entire source for migration. If desired, click the **Replication Rules** heading and expand the volumes under **Folders**. You will see that Carbonite Migrate automatically excludes particular files that cannot be used during the migration. If desired, you can exclude other files that you do not want to migrate, but be careful when excluding data. Excluded volumes, folders, and/or files may compromise the integrity of your installed applications. There are some volumes, folders, and files (identified in *italics* text) that you will be unable to exclude, because they are required for migration. For example, the boot files cannot be excluded because that is where the system state information is stored.

Volumes and folders with a green highlight are included completely. Volumes and folders highlighted in light yellow are included partially, with individual files or folders included. If there is no highlight, no part of the volume or folder is included. To modify the items selected, highlight a volume, folder, or file and click **Add Rule**. Specify if you want to **Include** or **Exclude** the item. Also, specify if you want the rule to be recursive, which indicates the rule should automatically be

applied to the subdirectories of the specified path. If you do not select **Recursive**, the rule will not be applied to subdirectories.

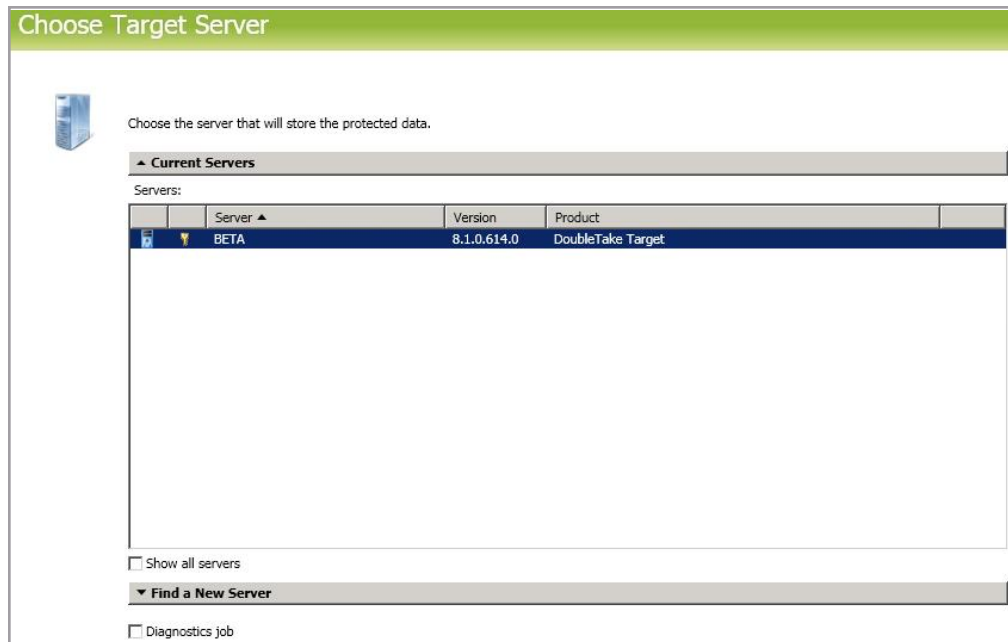
You can also enter wildcard rules, however you should do so carefully. Rules are applied to files that are closest in the directory tree to them. If you have rules that include multiple folders, an exclusion rule with a wild card will need to be added for each folder that it needs applied to. For example, if you want to exclude all .log files from D:\ and your rules include D:\, D:\Dir1 , and D:\Dir2, you would need to add the exclusion rule for the root and each subfolder rule. So you will need to add exclude rules for D:*.log , D:\Dir1*.log, and D:\Dir2*.log.

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



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

4. Click **Next** to continue.
5. Choose your target server. This is the Hyper-V server that will host the virtual machine that, after the migration, will become your new source.



- **Current Servers**—This list contains the servers currently available in your console session. Servers that are not licensed for the workflow you have selected and those not applicable to the workload type you have selected will be filtered out of the list. Select your target server from the list. If the server you are looking for is not displayed, enable **Show all servers**. The servers in red are not available for the source server or workload type you have selected. Hover your mouse over an unavailable server to see a reason why this

server is unavailable.

- **Find a New Server**—If the server you need is not in the **Current Servers** list, click the **Find a New Server** heading. From here, you can specify a server along with credentials for logging in to the server. If necessary, you can click **Browse** to select a server from a network drill-down list.



If you enter the target server's fully-qualified domain name, the Carbonite Replication 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.

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



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

-
7. You have many options available for your server migration job. Configure those options that are applicable to your environment.

Go to each page identified below to see the options available for that section of the **Set Options** page. After you have configured your options, continue with the next step on page 229.

- *General* on page 212
- *Replica Virtual Machine Location* on page 213
- *Replica Virtual Machine Configuration* on page 215
- *Replica Virtual Machine Volumes* on page 217
- *Replica Virtual Machine Network Settings* on page 219
- *Test Failover* on page 220
- *Failover Options* on page 222
- *Mirror, Verify & Orphaned Files* on page 223
- *Network Route* on page 226
- *Compression* on page 227
- *Bandwidth* on page 228

General

General

Job name:
alpha to beta

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

Replica Virtual Machine Location

Replica Virtual Machine Location

Local File Volume SMB Share

Select the volume and folder on the target server that will hold the replica virtual machine:

Volume	Total Size	Free Space	Owner
C:	100 GB	32.88 GB	
D:	308.37 GB	93.67 GB	
E:	136.02 GB	65.79 GB	

Full path where the replica virtual machine will be stored:

D:\HVRA_Replicas\alpha

Replica Virtual Machine Location

Local File Volume SMB Share

File server name: FileServer

Share name: ShareName

Folder: HVRA_Replicas\alpha

Full path where the replica virtual machine will be stored:

\\FileServer\ShareName\HVRA_Replicas\alpha



The options on this page will vary depending on the operating system of your target Hyper-V server.

- **Local File Volume**—Select this option if you want to store the new virtual server on a local volume on the target.
 - **Select the volume and folder on the target server that will hold the replica virtual machine**—Select one of the volumes from the list to indicate the volume on the target where you want to store the new virtual server when it is created. The target volume must have enough **Free Space** to store the source data.
 - **Full path where the replica virtual machine will be stored**—Specify a location on the selected **Volume** to store the replica of the source. By specifying an existing folder, you can reuse an existing virtual machine on your Hyper-V target created by a previous job.

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

Carbonite Migrate will skip the virtual disk creation steps when using a pre-existing disk, therefore, it is important that you place your pre-existing virtual disks in the location for **Full path where the replica virtual machine will be stored**.

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

- a. Create a job in a LAN environment, letting Carbonite Migrate create the virtual disk for you.
 - b. Complete the mirror process locally.
 - c. Delete the job and when prompted, do not delete the replica.
 - d. From the Hyper-V Manager, delete the replica virtual machine, which will delete the virtual machine configuration but will keep the associated hard disk files.
 - e. Shut down and move the Hyper-V target server to your remote site.
 - f. After the Hyper-V target server is back online at the remote site, create a new job for the same source server. Carbonite Migrate will reuse the existing hard disk files and perform a difference mirror over the WAN to bring the virtual machine up-to-date.
- **SMB Share**—If your target is Windows 2012, you can select this option if you want to store the new virtual server on an SMB 3.0 share. The file server where the share is located should not be the same server as your target Hyper-V server.
 - **File server name**—Specify the name of the file server where your SMB share is located.
 - **Share name**—Specify the share name on your file server where you want to store the new virtual server. Your target Hyper-V server must have permission to write to the share.
 - **Folder**—Specify a folder on the share where you want to store the new virtual server.

Replica Virtual Machine Configuration

	Source	Replica
Processors	1	1
Memory (MB)	4096	4096

Source Network Adapter	Replica Virtual Switch	Set VLAN on replica	Replica VLAN
Local Area Connection	Broadcom BCM5708C NetXtreme II GigE (N)	<input type="checkbox"/>	0



The fields shown in this section will vary depending on the operating system of your source and target servers.

- **Display name**—Specify the name of the replica virtual machine. This will be the display name of the virtual machine on the host system.
- **Hardware configuration**—Specify how you want the replica virtual machine to be created.
 - **Processors**—Specify how many processors to create on the new virtual machine. The number of processors on the source is displayed to guide you in making an appropriate selection. If you select fewer processors than the source, your clients may be impacted by slower responses.
 - **Memory**—Specify the amount of memory, in MB, to create on the new virtual machine. The memory on the source is displayed to guide you in making an appropriate selection. If you select less memory than the source, your clients may be impacted by slower responses.
- **Replica virtual machine generation**—If your source is running Windows 2012 or later and your target Hyper-V server is Windows 2012 R2 or later, you will have the choice to specify if the replica virtual machine is generation 1 or 2. If you do not have one of these configurations, you will not see this option.
- **Network adapter type**—Depending on your operating system and if you have selected a generation 1 virtual machine, you may be able to select the type of adapter, **Legacy** or **Synthetic**, to use on the replica virtual machine. This selection will apply to all adapters on the replica.
 - **Windows 2008**—Both legacy and synthetic adapters are supported by the operating system and Carbonite Migrate, so you will see this option.
 - **Windows 2012 or later**—Both legacy and synthetic adapters are supported by the operating system and Carbonite Migrate with one exception. If the source is

Windows 2012 or later and the target is Windows 2012 R2 or later, only the synthetic adapter is available for the operating system and Carbonite Migrate. You will see this option except for the one exception listed.

There is a limit of four legacy adapters and eight synthetic adapters.

- **Virtual switches**—Identify how you want to handle the network mapping after cutover. The **Source Network Adpater** column lists the NICs from the source. Map each one to a **Replica Virtual Switch**, which is a virtual network on the target. You can also choose to discard the source's NIC and IP addresses, or you can to failover the NIC and IP addresses but leave them in a not connected state.
- **Set VLAN on replica and Replica VLAN**—Enable the checkbox if you want to specify the VLAN ID. Specify the VLAN ID that will be used on the network adapter on the replica virtual machine.
- **Power on replica after failover**—By default, the replica virtual machine will automatically be powered on after the cutover process is complete. If you want the replica virtual machine to remain powered off, disable this option.

Replica Virtual Machine Volumes

Volume	Disk Size	Used Space	Replica Disk Size	Replica Disk Format	Storage Controller	Target Volume
C:	126.9 GB	10.71 GB	126.9 GB	Dynamic	IDE	D:

- **Replica Disk Size**—For each volume you are protecting, specify the size of the replica disk on the target. Be sure and include the value in MB or GB for the disk. The value must be at least the size of the specified **Used Space** on that volume.



In some cases, the replica virtual machine may use more virtual disk space than the size of the source volume due to differences in how the virtual disk's block size is formatted. To avoid this issue, make sure the replica can accommodate not just the size of all of the files but the size on disk as well.

- **Replica Disk Format**—For each volume you are protecting, specify the format of the disk, **Dynamic** or **Fixed**, that will be created on the replica virtual machine. Any disk format specification will be discarded if you are reusing a disk from the **Full path where the replica virtual machine will be stored** from the **Replica Virtual Machine Location** section.
- **Storage Controller**—For each volume you are protecting, specify the type of **Storage Controller** that you want to use for each volume on the target, keeping in mind the following guidelines.
 - If the protected virtual machine is less than Windows Server 2012 and/or the target Hyper-V version is less than Windows Server 2012 R2, Carbonite Migrate will always put the system volume and the first two data disks on an IDE controller. All subsequent volumes will be SCSI. If your target is less than Windows 2012, you will need to install the Hyper-V Integration Components to acquire a SCSI device after failover to attach the SCSI volumes to the replica virtual machine. See your Microsoft documentation for more information.
 - If the protected virtual machine is Windows 2012 or later and the target Hyper-V is Windows Server 2012 R2, you will have the choice of virtual machine generation.
 - **Generation 1**—If you selected a generation 1 virtual machine, Carbonite Migrate will always put the system volume on an IDE controller, but you can select the type of storage controller for all subsequent volumes. If you select SCSI controller, you will need Hyper-V Integration Components to acquire a SCSI device after failover in order to attach the SCSI volumes to the replica virtual machine. See your Microsoft documentation for more information on the Integration Components.
 - **Generation 2**—If you selected a generation 2 virtual machine, all volumes will be SCSI.
- **Target Volume**—For each volume you are protecting, specify the volume on the target where you want to store the virtual disk files for the new replica virtual machine. You can

select standalone volumes, CSVs, or cluster storage, depending on your target configuration. (Specify the location of the virtual machine configuration files under **Replica Virtual Machine Location**.)

Replica Virtual Machine Network Settings

Use advanced settings for replica virtual machine network configuration.

Network adapters:
Local Area Connection (112.42.74.29)

Source IP addresses:

IP Address	Subnet Mask
112.42.74.29	255.255.0.0

Replica IP addresses:

IP Address	Subnet Mask
112.52.74.29	255.255.0.0

Source Default Gateways:

112.42.48.9

Replica Default Gateways:

112.52.48.9

Source DNS Server addresses:

112.42.48.20

Replica DNS Server addresses:

112.52.48.20

- **Use advanced settings for replica virtual machine network configuration**—Select this option to enable the replica virtual machine network setting configuration. This setting is primarily used for WAN support.
- **Network adapters**—Select a network adapter from the source and specify the **Replica IP addresses**, **Replica Default Gateways**, and **Replica DNS Server addresses** to be used after cutover. 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 cutover 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 cutover.

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.

Test Failover

These options allow you to perform a test failover. Keep in mind the following for using test failover.

- The source, target, and protection job will remain online and uninterrupted during the test.
- The test will be performed using the test failover settings configured during job creation.
- The test will use the current data on the target.
- The test failover will take a snapshot of the current data on the target and create a new set of virtual disks. The data from the snapshot will be mirrored to the new set of disks using the same mirroring options as the protection job.
- Once the mirror is complete the replica virtual machine is automatically brought online using the new set of disks.
- The replica virtual machine will use the network settings specified in the test failover settings of the protection job.
- When you are finished with your test, undo it.
- When you undo a test failover, the new set of disks will be maintained or deleted as specified in the test failover settings of the protection job.
- At any time during a test failover, you can undo the test, perform a live failover, or failover to a snapshot. (Performing a live failover or failing over to a snapshot will automatically undo any test in progress.)

Test Failover

Network

Do not connect replica network adapters on test failover
 Connect and map replica network adapters on test failover

Map source virtual switches to target virtual switches for test failover:

Source Network Adapter	Replica Virtual Switch	Set VLAN on Replica	Replica VLAN
Local Area Connection	Not Connected	<input type="checkbox"/>	0

Configuration

Configure the volumes for the test replica virtual machine.

Volume	Replica Disk Format	Target Volume
C:	Dynamic	C:
E:	Fixed	E:

Delete test failover virtual disks

- **Do not connect replica network adapters on test failover**—Select this option if you do not want the replica virtual machine used for the test to be connected to the network.
- **Connect and map replica network adapters on test failover**—Select this option if you want the replica virtual machine used for the test to be connected to the network. You will need to map each **Source Network Adapter** to a **Target Virtual Switch** for the test. You can also choose to discard the source's NIC and IP addresses, or you can choose to failover the NIC and IP addresses but leave them in a not connected state. If you want to specify the VLAN ID, enable **Set VLAN on Replica** and specify the **Replica VLAN**.

- **Replica Disk Format**—For each volume you are protecting, specify the format of the disk, **Dynamic** or **Fixed**, that will be created on the test server.
 - **Target Volume**—For each volume you are protecting, specify the volume on the target where you want to store the virtual disk files for the test server. You can select standalone volumes, CSVs, or cluster storage, depending on your target configuration.
 - **Delete test failover virtual disks**—Select this option if you want to delete the new virtual disks created during the test failover process. If you disable this option, the new disks will not be deleted when you perform undo failover.
-



Be careful if you choose to connect the network adapters for a test cutover. Depending on your network adapter mappings, users may be able to access the target. Also, since the source is still online, there is a chance users may split between accessing the source or target.

Failover Options



Failover Options

- Wait for user to initiate failover
- Shutdown source server

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

Mirror, Verify & Orphaned Files

Mirror, Verify & Orphaned Files

Mirror Options

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

Compare file attributes. Send the attributes and bytes that differ.

Verification Options

Enable scheduled verification

Verify on this interval: 1 Days

Begin immediately

Begin at this time: 3/23/2017 10:53:32 AM

Report and comparison options

Report only

Report and mirror files

Compare file attributes and data

General Options

Calculate size of protected data upon connection

Delete orphaned files

- **Mirror Options**—Choose a comparison method and whether to mirror the entire file or only the bytes that differ in each file.
 - **Do not compare files. Send the entire file.**—Carbonite Migrate 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 option requires no time for comparison, but it can be slower because it sends the entire file. However, it is useful for configurations that have large data sets with millions of small files that are frequently changing and it is more efficient to send the entire file. You may also need to use this option if configuration management policies require sending the entire file.
 - **Compare file attributes. Send the attributes and bytes that differ.**—Carbonite Migrate will compare file attributes and will mirror only the attributes and bytes that are different. This option is the fastest comparison method and fastest mirror option. Files that have not changed can be easily skipped. Also files that are open and require a checksum mirror can be compared.
 - **Compare file attributes and data. Send the attributes and bytes that differ.**—Carbonite Migrate will compare file attributes and the file data and will mirror only the attributes and bytes that are different. This comparison method is not as fast because every file is compared, regardless of whether the file has changed or is open. However, sending only the attributes and bytes that differ is the fastest mirror 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 Carbonite Migrate



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, Carbonite Migrate 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 Carbonite Migrate 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.


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

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

If you want to move orphaned files rather than delete them, you can configure this option along with the move deleted files feature to move your orphaned files to the specified deleted files directory. See *Target server properties* on page 47 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 Carbonite Migrate 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 Carbonite Migrate log, and during difference mirrors, verifications, and restorations, all orphaned file processing messages are logged to the Carbonite Migrate log. The orphaned file log is located in the **Logging folder** specified for the source. See *Log file properties* on page 52 for details on the location of that folder. The orphaned log file is appended to during each orphaned file processing during a mirror, and the log file will be a maximum of 50 MB.

Network Route

 **Network Route**

Send data to the target server using this route:




This section is not applicable to clustered environments.

By default, Carbonite Migrate 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 Carbonite Migrate traffic. For example, you can separate regular network traffic and Carbonite Migrate traffic on a machine with multiple IP addresses. You can also select or manually enter a public IP address (which is the public IP address of the server's router) if you are using a NAT environment. The **Management Service port** and **Replication Service port** can be disregarded. It is used for other job types.



The IP address used on the source will be determined through the Windows route table.

Compression



To help reduce the amount of bandwidth needed to transmit Carbonite Migrate 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.

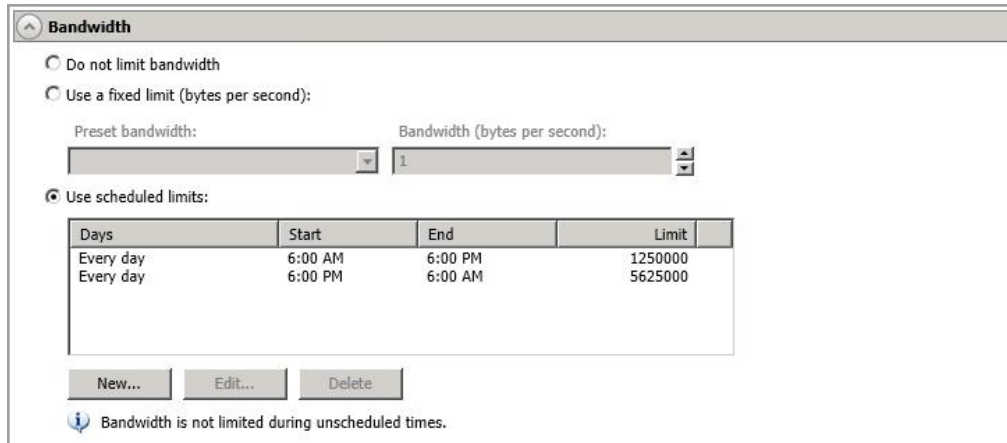
Keep in mind that the process of compressing data impacts processor usage on the source. If you notice an impact on performance while compression is enabled in your environment, either adjust to a lower level of compression, or leave compression disabled. Use the following guidelines to determine whether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



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

Bandwidth



Bandwidth


Do not limit bandwidth

Use a fixed limit (bytes per second):

Preset bandwidth: Bandwidth (bytes per second):

Use scheduled limits:

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

 Bandwidth is not limited during unscheduled times.

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



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

- **Do not limit bandwidth**—Carbonite Migrate will transmit data using 100% bandwidth availability.
- **Use a fixed limit**—Carbonite Migrate 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**—Carbonite Migrate 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**.

8. Click **Next** to continue.
9. Carbonite Migrate 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 Carbonite Migrate correct the problem for you. For those errors that Carbonite Migrate 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 *Carbonite Availability and Carbonite Migrate Reference Guide* for details on the various Carbonite Migrate log files.

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



Jobs in a NAT environment may take longer to start.

Once a job is created, do not change the name of underlying hardware components used in the job. For example, volume and datastore names or network adapter and virtual switch names. Any component used by name in your job must continue to use that name throughout the lifetime of job. If you must change a name, you will need to delete the job and re-create it using the new component name.

Managing and controlling full server to Hyper-V migration jobs

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

The jobs displayed in the right pane depend on the server group folder selected in the left pane. Every job for each server in your console session is displayed when the **Jobs on All Servers** group is selected. If you have created and populated server groups (see *Managing servers* on page 18), then only the jobs associated with the server or target servers in that server group will be displayed in the right pane.

- *Overview job information displayed in the top right pane* on page 231
- *Detailed job information displayed in the bottom right pane* on page 234
- *Job controls* on page 236


Overview job information displayed in the top right pane


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


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


Column 1 (Blank)

The first blank column indicates the state of the job.

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

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

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

 The job is in an unknown state.

Job

The name of the job

Source Server

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

Target Server

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

Job Type

Each job type has a unique job type name. This job is a Full Server to Hyper-V Migration job. For a complete list of all job type names, press F1 to view the Carbonite Replication Console online help.

Activity

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

Mirror Status

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

Replication Status

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

Transmit Mode

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

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

Operating System

The job type operating system

Detailed job information displayed in the bottom right pane

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

Name

The name of the job

Target data state

- **OK**—The data on the target is in a good state.
- **Mirroring**—The target is in the middle of a mirror process. The data will not be in a good state until the mirror is complete.
- **Mirror Required**—The data on the target is not in a good state because a remirror is required. This may be caused by an incomplete or stopped mirror or an operation may have been dropped on the target.
- **Busy**—The source is low on memory causing a delay in getting the state of the data on the target.
- **Not Loaded**—Carbonite Migrate 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 38.

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 38.

Replication queue

The total number of replication bytes in the source queue

Disk queue

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

Recovery point latency

The length of time replication is behind on the target compared to the source. This is the time period of replication data that would be lost if a failure were to occur at the current time. This value represents replication data only and does not include mirroring data. If you are mirroring and failover, the data on the target will be at least as far behind as the recovery point latency. It could potentially be further behind depending on the

circumstances of the mirror. If mirroring is idle and you failover, the data will only be as far behind as the recovery point latency time.

Bytes sent

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

Bytes sent (compressed)

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

Connected since

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

Recent activity

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

Additional information

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

Job controls

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

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

View Job Details

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

Edit Job Properties

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

Delete

Stops (if running) and deletes the selected jobs.

Provide Credentials

Changes the login credentials that the job (which is on the target machine) uses to authenticate to the servers in the job. This button opens the Provide Credentials dialog box where you can specify the new account information and which servers you want to update. See *Providing server credentials* on page 31. You will remain on the **Jobs** page after updating the server credentials. If your servers use the same credentials, make sure you also update the credentials on the **Servers** page so that the Carbonite Replication Console can authenticate to the servers in the console session. See *Managing servers* on page 18.

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 Carbonite Migrate queue was not exhausted during

the time the job was paused. If the Carbonite Migrate queue was exhausted during the time the job was paused, the job will restart mirroring and replication.

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

Pause

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

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

Stop

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

Take Snapshot

Snapshots are not applicable to migration jobs.

Manage Snapshots

Snapshots are not applicable to migration jobs.

Failover or Cutover

Starts the cutover process. See *Cutting over full server to Hyper-V migration jobs* on page 250 for the process and details of cutting over a full server to Hyper-V migration job.

Failback

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

Restore

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

Reverse

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

Undo Failover or Cutover

Cancels a test cutover by undoing it. This resets the servers and the job back to their original state. See *Cutting over full server to Hyper-V migration jobs* on page 250 for the process and details of undoing a cut over full server to Hyper-V migration 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, Carbonite Migrate 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, Carbonite Migrate 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.**—Carbonite Migrate 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 option requires no time for comparison, but it can be slower because it sends the entire file. However, it is useful for configurations that have large data sets with millions of small files that are frequently changing and it is more efficient to send the entire file. You may also need to use this option if configuration management policies require sending the entire file.
 - **Compare file attributes. Send the attributes and bytes that differ.**—Carbonite Migrate will compare file attributes and will mirror only the attributes and bytes that are different. This option is the fastest comparison method and fastest mirror option. Files that have not changed can be easily skipped. Also files that are open and require a checksum mirror can be compared.
 - **Compare file attributes and data. Send the attributes and bytes that differ.**—Carbonite Migrate will compare file attributes and the file data and will mirror only the attributes and bytes that are different. This comparison method is not as fast because every file is compared, regardless of whether the file has changed or is open. However, sending only the attributes and bytes that differ is the fastest mirror option.
- **Calculate size of protected data before mirroring**—Specify if you want

Carbonite Migrate 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**—Carbonite Migrate will transmit data using 100% bandwidth availability.
 - **Fixed bandwidth limit**—Carbonite Migrate 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 Carbonite Migrate 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 Carbonite Migrate processing, not the entire server.

While the target is paused, the Carbonite Migrate 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, Carbonite Migrate will automatically disconnect the connections and attempt to reconnect them.

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

Filter

Select a filter option from the drop-down list to only display certain jobs. You can display **Healthy jobs**, **Jobs with warnings**, or **Jobs with errors**. To clear the filter, select **All jobs**. If you have created and populated server groups, then the filter will only apply to the jobs associated with the server or target servers in that server group. See *Managing servers* on page 18.

Search

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

Overflow Chevron

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

Viewing full server to Hyper-V migration job details

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

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





Job name

The name of the job

Job type

Each job type has a unique job type name. This job is a Full Server to Hyper-V Migration job. For a complete list of all job type names, press F1 to view the Carbonite Replication Console online help.

Health

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

Activity

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

Connection ID

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

Transmit mode

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

Target data state

- **OK**—The data on the target is in a good state.
- **Mirroring**—The target is in the middle of a mirror process. The data will not be in a good state until the mirror is complete.
- **Mirror Required**—The data on the target is not in a good state because a remirror is required. This may be caused by an incomplete or stopped mirror or an operation may have been dropped on the target.
- **Busy**—The source is low on memory causing a delay in getting the state of the data on the target.
- **Not Loaded**—Carbonite Migrate 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 Carbonite Migrate transmissions.

Compression

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

Encryption

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

Bandwidth limit

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

Connected since

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

Additional information

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

Mirror status

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

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

Mirror percent complete

The percentage of the mirror that has been completed

Mirror remaining

The total number of mirror bytes that are remaining to be sent from the source to the target. This value may be zero if you have enabled **Mirror only changed files when source reboots** on the *Server setup properties* on page 38.

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 38.

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 Carbonite Migrate driver. Check the Event Viewer for driver related issues.
- **Unknown**—The console cannot determine the status.

Replication queue

The total number of replication bytes in the source queue

Disk queue

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

Bytes sent

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

Bytes sent compressed

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

Recovery point latency

The length of time replication is behind on the target compared to the source. This is the time period of replication data that would be lost if a failure were to occur at the current time. This value represents replication data only and does not include mirroring data. If you are mirroring and failover, the data on the target will be at least as far behind as the recovery point latency. It could potentially be further behind depending on the circumstances of the mirror. If mirroring is idle and you failover, the data will only be as far behind as the recovery point latency time.

Mirror start time

The UTC time when mirroring started

Mirror end time

The UTC time when mirroring ended

Total time for last mirror

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

Validating a full server to Hyper-V migration job

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

1. From the **Jobs** page, highlight the job and click **View Job Details** in the toolbar.
2. In the **Tasks** area on the right on the **View Job Details** page, click **Validate job properties**.
3. Carbonite Migrate 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 Carbonite Migrate correct the problem for you. For those errors that Carbonite Migrate cannot correct automatically, you will need to modify the source or target to correct the error, or you can select a different target. You must revalidate the selected servers, by clicking **Recheck**, until the validation check passes without errors.

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

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

Editing a full server to Hyper-V migration job

Use these instructions to edit a full server to Hyper-V migration job.

1. From the **Jobs** page, highlight the job and click **Edit Job Properties** in the toolbar. (You will not be able to edit a job if you have removed the source of that job from your Carbonite Replication Console session or if you only have Carbonite Migrate monitor security access.)
2. You will see the same options available for your full server to Hyper-V migration job as when you created the job, but you will not be able to edit all of them. If desired, edit those options that are configurable for an existing job. See *Creating a full server to Hyper-V migration job* on page 209 for details on each job option.



Changing some options may require Carbonite Migrate to automatically disconnect, reconnect, and remirror the job.

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

-
3. If you want to modify the workload items or replication rules for the job, click **Edit workload or replication rules**. Modify the **Workload item** you are protecting, if desired. Additionally, you can modify the specific **Replication Rules** for your job.

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

You can also enter wildcard rules, however you should do so carefully. Rules are applied to files that are closest in the directory tree to them. If you have rules that include multiple folders, an exclusion rule with a wild card will need to be added for each folder that it needs applied to. For example, if you want to exclude all .log files from D:\ and your rules include D:\, D:\Dir1 , and D:\Dir2, you would need to add the exclusion rule for the root and each subfolder rule. So you will need to add exclude rules for D:*.log , D:\Dir1*.log, and D:\Dir2*.log.

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

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



If you remove data from your workload and that data has already been sent to the target, you will need to manually remove that data from the target. Because the data you

removed is no longer included in the replication rules, Carbonite Migrate orphan file detection cannot remove the data for you. Therefore, you have to remove it manually.

4. Click **Next** to continue.
5. Carbonite Migrate 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 Carbonite Migrate correct the problem for you. For those errors that Carbonite Migrate 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.

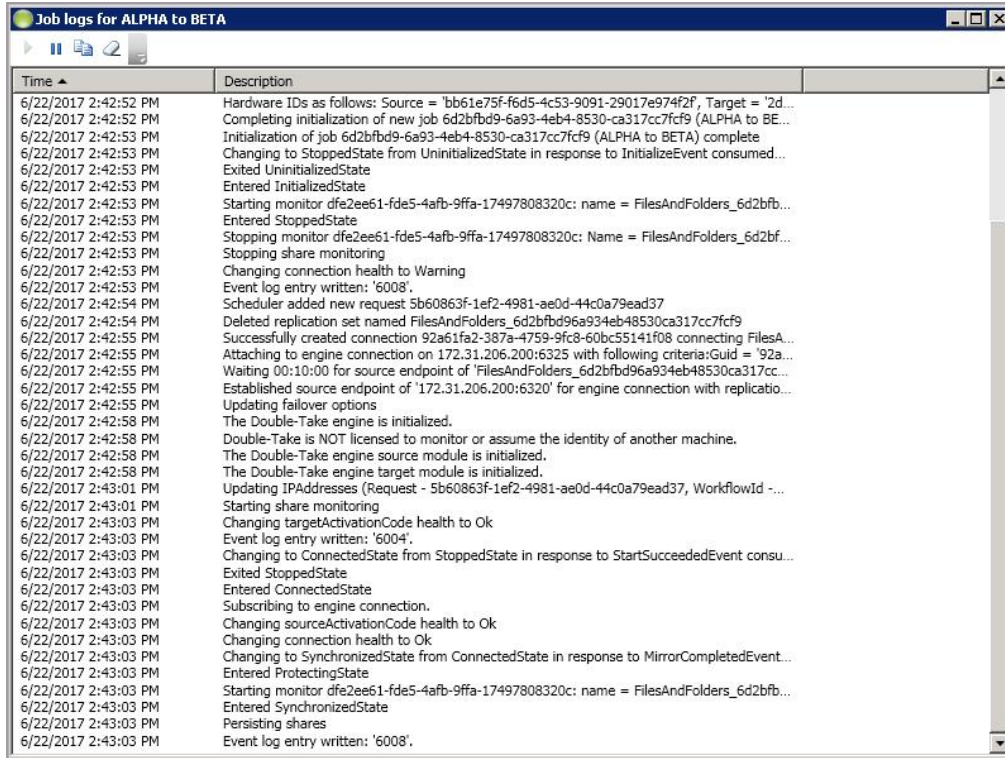
If you receive a path transformation error during job validation indicating a volume does not exist on the target server, even though there is no corresponding data being protected on the source, you will need to manually modify your replication rules. Go back to the **Choose Data** page and under the **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 *Carbonite Availability and Carbonite Migrate Reference Guide* for details on the various Carbonite Migrate log files.

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

Viewing a full server to Hyper-V migration job log

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



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

Start

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

Pause

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

Copy

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

Clear

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

Time

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

Description

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

Cutting over full server to Hyper-V migration jobs

When the migration mirror has completed, the target may or may not reboot automatically depending on your selection for **Wait for user intervention before cutover**. If you disabled user intervention, the target will reboot automatically to complete the migration process. If you enabled user intervention, when the migration mirror is complete, the status will change to **Protecting**. Use this time to complete any necessary tasks. When you are ready to complete the migration, use the following instructions to cutover.

1. On the **Jobs** page, highlight the job that you want to cutover and click **Failover, Cutover, or Recover** in the toolbar.
2. Select the type of cutover to perform.
 - **Cutover to live data**—Select this option to initiate a full, live cutover using the current data on the target. The source may be automatically shut down if it is still running, depending on your job configuration. The protection job is stopped and the replica virtual machine is started on the target with full network connectivity.
 - **Perform test cutover**—Select this option to perform a test cutover.
 - The source, target, and protection job will remain online and uninterrupted during the test.
 - The test will be performed using the test failover settings configured during job creation.
 - The test will use the current data on the target.
 - The test failover will take a snapshot of the current data on the target and create a new set of virtual disks. The data from the snapshot will be mirrored to the new set of disks using the same mirroring options as the protection job.
 - Once the mirror is complete the replica virtual machine is automatically brought online using the new set of disks.
 - The replica virtual machine will use the network settings specified in the test failover settings of the protection job.
 - When you are finished with your test, undo it.
 - When you undo a test failover, the new set of disks will be maintained or deleted as specified in the test failover settings of the protection job.
 - At any time during a test failover, you can undo the test, perform a live failover, or failover to a snapshot. (Performing a live failover or failing over to a snapshot will automatically undo any test in progress.)
 - **Cutover to a snapshot**—This option is not available for migration jobs.
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 cutover begins. The advantage to this option is that all of the data that the target has received will be applied before cutover 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 cutover will begin immediately. The advantage to this option is that cutover will occur immediately. The disadvantage is that any data in the target queue will be lost.

4. When you are ready to begin cutover, click **Cutover**.



Once cutover has started, do not reboot the target. If the cutover process is interrupted, it may fail.



If your source was disconnected from the network during the cutover process, the original job that was mirroring the data will still be active on that machine. Do not bring that machine back on the network. If you do, the original job will attempt to begin mirroring data again which could result in data loss. While the machine is still disconnected from the network, stop the Double-Take service, delete the file connect.sts located in the Carbonite Migrate installation directory, restart the Double-Take service, and bring the machine on the network if desired.

Because the Windows product activation is dependent on hardware, you may need to reactivate your Windows registration after cutover. The reactivation depends on several factors including service pack level, Windows edition, and your licensing type. If a target comes online after cutover with an activation failure, use the steps below appropriate for your license type. Additionally, if you are using Windows 2012, you may only have 60 minutes to complete the reactivation process until Windows activation tampering automatically shuts down your server.

- **Retail licensing**—Retail licensing allows the activation of a single operating system installation.
 1. Open the **System** applet in Windows **Control Panel**.
 2. Under **Windows activation** at the bottom of the page, click **Change product key**.
 3. Enter your retail license key. You may need access to the Internet or to call Microsoft to complete the activation.
- **MAK volume licensing**—Multiple Activation Key (MAK) licensing allows the activation of multiple operating system installations using the same activation key.
 1. View or download the Microsoft Volume Activation Deployment Guide from the Microsoft web site.
 2. Using an administrative user account, open a command prompt and follow the instructions from the deployment guide to activate MAK clients. Multiple reboots may be necessary before you can access a command prompt. You may need access to the Internet or to call Microsoft to complete the activation.
- **KMS volume licensing**—Key Management Service (KMS) licensing allows IT professionals to complete activations on their local network without contacting Microsoft.
 1. View or download the Microsoft Volume Activation Deployment Guide from the Microsoft web site.
 2. Using an administrative user account, open a command prompt and follow the instructions from the deployment guide to convert a MAK activation client to a KMS client. Multiple reboots may be necessary before you can access a command prompt.

If your source is running on Windows Server 2008 and the target replica has one or more SCSI drives, then after cutover the CD/DVD ROM will not be allocated. If the CD/DVD ROM is required, you will need to edit the virtual machine settings to add a CD/DVD ROM after cutover. By not allocating a CD/DVD ROM under these specific conditions, drive letter consistency will be guaranteed.

If your source is a virtual machine running on an ESX host and you have VMware Tools installed, VMware Tools will be disabled during cutover.

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

After failover is complete, an additional NIC may be seen when looking at hidden devices or a system scan. These NICs are not loaded so they are not consuming any resources and are not active. You can safely disregard these hidden NICs.

5. If you performed a test cutover, you can undo it by selecting **Undo Failover or Cutover** in the toolbar. The replica virtual machine on the target will be shut down and, if configured, the virtual disks used for the test cutover will be deleted.

Chapter 9 Simulating migration

Carbonite Migrate offers a simple way for you to simulate migration 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 Carbonite Migrate operations. Use the Carbonite Replication Console to automatically collect this data. It gathers Carbonite Migrate log files; Carbonite Migrate and system settings; network configuration information such as IP, WINS and DNS addresses; and other data which may be necessary in evaluating Carbonite Migrate performance.

1. From the Carbonite Replication Console, on the **Servers** page, right-click the source where you will be running the TDU, select **Gather Support Diagnostics**, and specify a location to store the zipped diagnostics information. It may take several minutes for the diagnostics to finish processing. After it is complete, a .zip file containing the information gathered will be created. The file name is based on the machine name.
2. Establish a migration 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 Carbonite Migrate statistics file. See the *Carbonite Availability and Carbonite Migrate Reference Guide* for details on DTStat.
4. After your simulation is complete, repeat step 1 to gather diagnostic again.

Chapter 10 Special network configurations

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

- *Firewalls* on page 255
- *IP and port forwarding* on page 256
- *Macintosh shares* on page 259
- *NFS Shares* on page 260

Firewalls

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

- **Carbonite Migrate ports**—Ports 6320, 6325, and 6326 are used for Carbonite Migrate communications and must be open on your firewall. Open UDP and TCP for both inbound and outbound traffic.
- **Microsoft WMI and RPC ports**—Some features of Carbonite Migrate and the Carbonite Replication Console use WMI (Windows Management Instrumentation) which uses RPC (Remote Procedure Call). By default, RPC will use ports at random above 1024, and these ports must be open on your firewall. RPC ports can be configured to a specific range by specific registry changes and a reboot. See the [Microsoft Knowledge Base article 154596](#) for instructions.
- **Microsoft File Share and Directory ports**—Carbonite Migrate push installations will also rely on File Share and Directory ports, which must be open on your firewall. Check your Microsoft documentation if you need to modify these ports.
 - Microsoft File Share uses ports 135 through 139 for TCP and UDP communications.
 - Microsoft Directory uses port 445 for TCP and UDP communications.
- **ESX ports**—If you are using VirtualCenter or an ESX host, port 443 is also required and must be opened.

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

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

IP and port forwarding

As outlined in the requirements, Carbonite Migrate supports IP and port forwarding in NAT environments with the following caveats.

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

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

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

Add Servers

Identify the servers in your environment that you want to manage. The servers you add here appear on the Servers page.

Manual Entry | Automatic Discovery

Server: 112.47.12.7

User name: domain\administrator

Password:

Domain:

Management Service port: 1025 Use default port

Add

Servers to be added:

Server	Details
--------	---------


Remove Remove All

OK Cancel

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



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



- For full server jobs, you can enable an option to change the target port to match the source port during cutover. This option is in the job creation wizard on the **Set Options** page in the **Failover Options** section.

Failover Options

- Wait for user to initiate failover
- Shutdown source server
- Change target ports to match source during cutover

Target scripts

Pre-failover script: ... Arguments:

Delay failover until script completes

Post-failover script: ... Arguments:

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

Network Route

Send data to the target server using this route: Management Service port: Use default port

Replication Service port: Use default port

Receive requests from the target server using this route: Use default route

Reverse Protection and Routing

Send data to the target server using this route: Management Service port: Use default port

Replication Service port: Use default port

Receive requests from the target server using this route: Use default route

Enable reverse protection

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

Select a reserved IP address on the source:

Send data to source after reverse using this route: Replication Service port: Use default port

Select a reserved IP address on the target:

Macintosh shares

A share is any volume, drive, or directory resource that is shared across a network. During failover, the target can assume or add any source shares so that they remain accessible to the end users. Automatic share failover only occurs for standard Windows file system shares. Other shares, including Macintosh volumes, must be configured for failover through the failover scripts or created manually on the target.

1. On your target, set the File Server for Macintosh service to manual startup. This allows the post-failover script on the target to control when the service starts on the target.
2. Create each volume on the target machine exactly as it exists on the source. Use the Shared Folder wizard to configure each volume as a Macintosh-accessible volume. Follow these steps to start the wizard.
 - a. Open the Control Panel and click **Administrative Tools**.
 - b. Select **Configure Your Server**.
 - c. In the Configure Your Server window, click the **File Server** link.
 - d. Click **Start the Shared Folder wizard** to start the wizard, and then follow the directions provided by the wizard. On the Create Shared Folders screen, you must enable **Apple Macintosh**.



You can automate the creation of the volumes during the failover process by using the macfile volume command in the post-failover batch file. For detailed information on how to use this command, see your Windows reference guide.

3. On the target machine, copy the chngname utility, chngname.exe, from the \tools\Win2K directory of the Carbonite Migrate DVD or from the Carbonite support web site to the directory where Carbonite Migrate is installed.
4. Add the following to your failover script.

```
rem Commands for Macintosh-accessible volume failover
rem The chngname utility (chngname.exe) must be located in the same
rem directory where Carbonite Migrate is installed.
rem The following command temporarily changes the name of the server. You
rem will need to replace <drive>:\<directory>\ with the location of
rem your Carbonite Migrate chngname utility and replace
rem source_name with the name of the source machine.
<drive>\<directory>\chngname /s source_name
rem The following command starts the File Server for Macintosh service
net start "File server for Macintosh"
rem The following command changes the name of the server back to its
rem original name. You will need to replace <drive>:\<directory>\ with
rem the location of your Carbonite Migrate chngname utility.
<drive>\<directory>\chngname /t
```

In the event of a failure, the Macintosh clients must remap the volume in order to access it. From the Macintosh client, use the Chooser to select the volume that needs to be remapped.

NFS Shares

A share is any volume, drive, or directory resource that is shared across a network. During failover, the target can assume or add any source shares so that they remain accessible to the end users. Automatic share failover only occurs for standard Windows file system shares. Other shares, including NFS shares, must be configured for failover through the failover scripts or created manually on the target.

1. On your target, set the NFS service to manual startup. This allows the post-failover script on the target to control when the service starts on the target.
2. Create each shared drive or directory on the target exactly as it exists on the source. Configure each drive or directory as an NFS share by following these steps.
 - a. Right-click the drive or directory that you want to share, select **Sharing**, and click the **NFS Sharing** tab on the Program Files Properties dialog box.
 - b. Enable **Share this folder**, provide the name of the share, and click **OK**.
3. On the target machine, copy the chngname utility, chngname.exe, from the \tools\Win2K directory of the Carbonite Migrate DVD or from the support web site to the directory where Carbonite Migrate is installed.
4. Add the following to your failover script.

```
rem Commands for NFS share failover
rem The chngname utility (chngname.exe) must be located in the same
rem directory where Carbonite Migrate is installed.
rem The following command temporarily changes the name of the server. You
rem will need to replace <drive>:\<directory>\ with the location of
rem your Carbonite Migrate chngname utility and replace
rem source_name with the name of the source machine.
<drive>\<directory>\chngname /s source_name
rem The following command starts the NFS service
net start "Server for NFS"
```

In the event of a failure, the clients must remount the shares in order to access them.