

Carbonite Availability for Windows

User's Guide



Notices

Carbonite Availability for Windows User's Guide, version 8.4.1, Wednesday, January 13, 2021

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.

© 2021 Carbonite, Inc. All rights reserved.

Contents

Chapter 1 Carbonite Availability overview	6
Core operations	7
Carbonite Availability workloads	10
Supported configurations	13
Chapter 2 Requirements	20
Mirroring and replication capabilities	21
Chapter 3 Carbonite Replication Console	27
Carbonite Replication Console requirements	29
Console options	30
Chapter 4 Managing servers	33
Adding servers	43
Providing server credentials	46
Viewing server details	47
Editing server properties	49
General server properties	50
Server licensing	51
Server setup properties	54
Carbonite Availability queue	57
Source server properties	61
Target server properties	63
E-mail notification configuration	65
Script credentials	67
Log file properties	68
Verification log	70
Viewing server events	72
Viewing server logs	73
Managing VMware servers	75
Managing snapshots	76
Snapshot states	77
Chapter 5 Files and folders protection	80
Files and folders requirements	81
Creating a files and folders job	87
Managing and controlling files and folders jobs	114
Viewing files and folders job details	124
Validating a files and folders job	128
Editing a files and folders job	129
Viewing a files and folders job log	131
Failing over files and folders jobs	133
Failback and restoration for files and folders jobs	134
Restoring then failing back files and folders jobs	135
Failing back then restoring files and folders jobs	137
Chapter 6 Full server protection	139
Full server requirements	140
Creating a full server job	149


Managing and controlling full server jobs	173
Viewing full server job details	183
Validating a full server job	187
Editing a full server job	188
Viewing a full server job log	190
Failing over full server jobs	192
Reversing full server jobs	196
Reversing full server jobs manually	198
Chapter 7 SQL protection	201
SQL requirements	202
Creating a SQL job	209
Managing and controlling SQL jobs	232
Viewing SQL job details	242
Validating a SQL job	246
Editing a SQL job	247
Viewing a SQL job log	249
Failing over SQL jobs	251
Restoring then failing back SQL jobs	253
Chapter 8 Full server to Hyper-V protection	254
Full server to Hyper-V requirements	255
Creating a full server to Hyper-V job	261
Managing and controlling full server to Hyper-V jobs	283
Viewing full server to Hyper-V job details	293
Validating a full server to Hyper-V job	297
Editing a full server to Hyper-V job	298
Viewing a full server to Hyper-V job log	300
Failing over full server to Hyper-V jobs	302
Reversing protection after failover for full server to Hyper-V jobs	306
Chapter 9 Full server to ESX protection	307
Full server to ESX requirements	308
Creating a full server to ESX job	315
Managing and controlling full server to ESX jobs	341
Viewing full server to ESX job details	351
Validating a full server to ESX job	355
Editing a full server to ESX job	356
Viewing a full server to ESX job log	358
Failing over full server to ESX jobs	360
Reversing protection after failover for full server to ESX jobs	364
Chapter 10 Simulating protection	365
Chapter 11 Special network configurations	366
Firewalls	367
IP and port forwarding	368
Domain controllers	371
NetBIOS	372
WINS	373
DNS	375
Non-Microsoft DNS	383

Macintosh shares	385
NFS Shares	386
Chapter 12 Recommended optimizations	387
Planning	388
Installation optimizations	389
General optimizations	390
Full server optimizations	394
Application optimizations	395
Chapter 13 Security	396
Adding users to the security groups	397
Changing the account used to run the Double-Take service on Windows servers ...	398

Chapter 1 Carbonite Availability overview

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

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



Core operations


Carbonite Availability performs three basic types of operations.

- *Mirroring* on page 7—The initial copy or subsequent resynchronization of selected data
- *Replication* on page 8—The on-going capture of byte-level file changes
- *Failover* on page 9—The ability to stand-in for a server, in the event of a failure

Mirroring

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

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




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

Replication

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

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



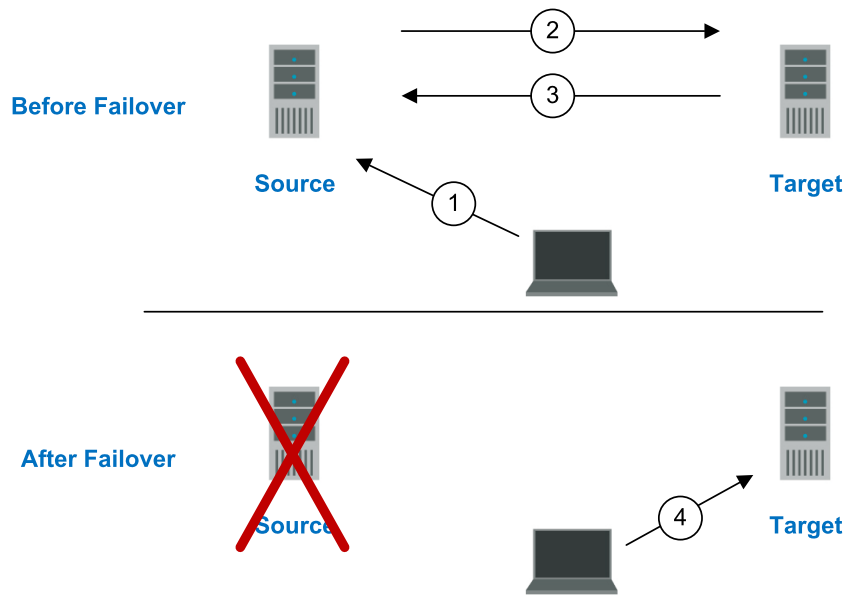
1. A user or application updates part of a file.
2. Only the changed portion of the file is replicated to the target.
3. An up-to-date copy of the file is maintained on the target.

Failover

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

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

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



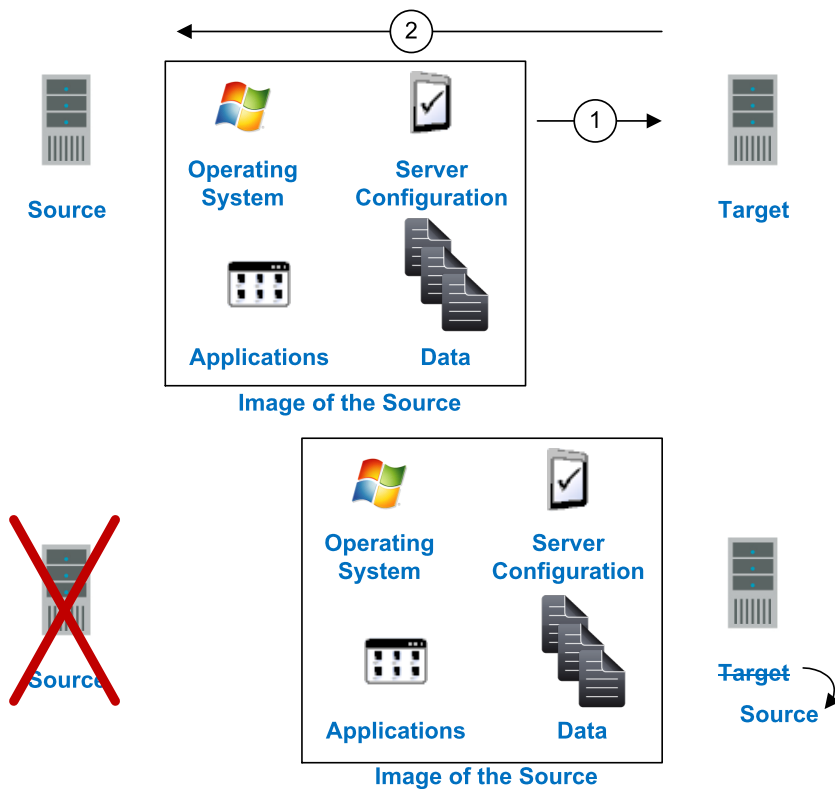
1. User and application requests are sent to the source name or IP address.
2. Data on the source is mirrored and replicated to the target.
3. The target monitors the source for failure.
4. In the event the source fails, the target stands in for the source. User and application requests are still sent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

Carbonite Availability workloads

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

Full server protection

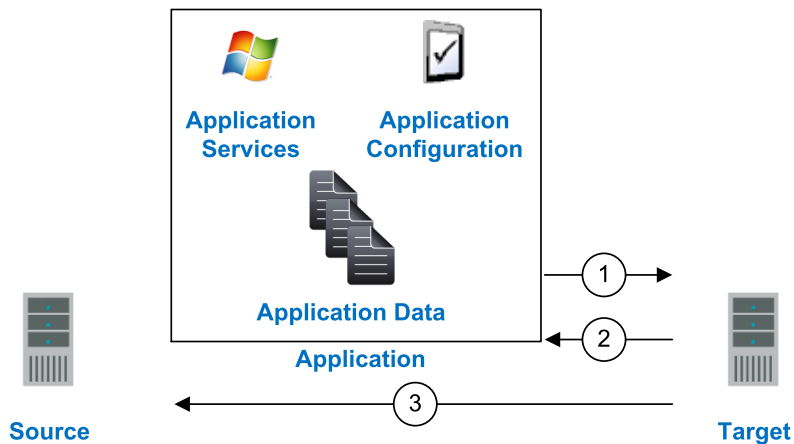
Full server protection provides high availability for an entire server, including the system state, which is the server's configured operating system and applications. You identify your source, which is the server you want to protect, and your target, which is the server that will stand-in for the source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



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

Application protection

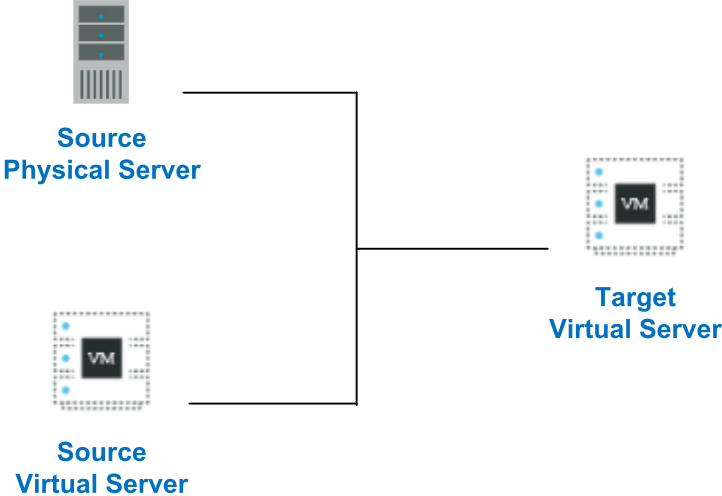
Application protection provides high availability for Microsoft SQL Server. You identify your source, which is the server running the application, and your target, which is the server that will stand-in for the source in the event the source fails. Carbonite Availability will gather information from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



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

Virtual protection

Virtual protection provides high availability from physical or virtual machines to virtual machines. You identify your source, which is the server you want to protect. Your source can be a physical server or a virtual machine. Your target is a Hyper-V or ESX server that will host a virtual machine that is a replica of the source. Carbonite Availability monitors the source for a failure. In the event of a source failure, the replica virtual machine on the target can stand-in allowing end-users to continue accessing data and/or applications.



Supported configurations


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



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

- *One to one, active/standby* on page 14
- *One to one, active/active* on page 15
- *Many to one* on page 16
- *One to many* on page 17
- *Chained* on page 18
- *Single server* on page 19

One to one, active/standby



Description

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

Applications

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

Considerations

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

One to one, active/active



Description

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


Applications

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

Considerations

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

Many to one



Description

Many source servers are protected by one target server.


Applications

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

Considerations

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

One to many



Description

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


Applications

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

Considerations

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

Chained



Description

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


Applications

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

Considerations

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

Single server



Description

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

Applications

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

Considerations

- This configuration does not provide high availability for the entire server.
- This configuration must be configured carefully so that an infinite loop is not created.
- This configuration should be limited to a single Carbonite Availability job.
- This configuration should be used sparingly. If possible, you should attach the target volumes to another server and use a one to one configuration.

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 requirements* on page 81
- *SQL requirements* on page 202
- *Full server requirements* on page 140
- *Full server to Hyper-V requirements* on page 255
- *Full server to ESX requirements* on page 308

Mirroring and replication capabilities

For Windows source servers, Carbonite Availability 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 Availability 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 Availability 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 Availability 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 Availability 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 Availability 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 and later.
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. Depending on your job type, you may not be able to use Windows 2016 ReFS to Windows 2019 ReFS or Windows 2019 ReFS to Windows 2016 ReFS. (ReFS is formatted differently on Windows 2016 and Windows 2019.) See the requirements for your job type for details. 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 Availability 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 Availability 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 Availability 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 Availability supports transactional NTFS (TxF) write operations, with the exception of TxF SavePoints (intermediate rollback points).
- a. With transactional NTFS and Carbonite Availability 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 Availability 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 Availability **Target Data State** will

indicate **Transactions Pending**. As the pending transactions are committed or aborted, Carbonite Availability 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 Availability log file for a list of files with pending transactions. As transactions are committed or aborted, the list is updated until all transactions are complete, and the **Target Data State** is **OK**.

- b. During replication, transactional operations will be processed on the target identically as they are on the source. If a transaction is committed on the source, it will be committed on the target. If a transaction is aborted on the source, it will be aborted on the target.
 - c. When failover occurs any pending transactions on the target will be aborted.
 - d. Carbonite Availability restore functions as a mirror, except the roles of the source and target are reversed. The transacted view of the data on the target is restored to the source. As pending transactions are committed or aborted on the target, Carbonite Availability restores any necessary changes to the source. Once all pending transactions are completed, the restoration is complete and replication will continue from the target to the source.
 - e. If you have restored your data before starting the failback process, make sure the restoration process does not have pending transactions and is complete before starting failback. If you are restoring your data after the failback the process has completed, users will not be accessing the data once failback occurs, so there are no opportunities for pending transactions.
11. Carbonite Availability 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 Availability 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 file 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 Availability 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 Availability 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 Availability 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 Availability 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 Availability. 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 Availability 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 server with a full server protection job. If it is not in its default installation directory, failover will fail.

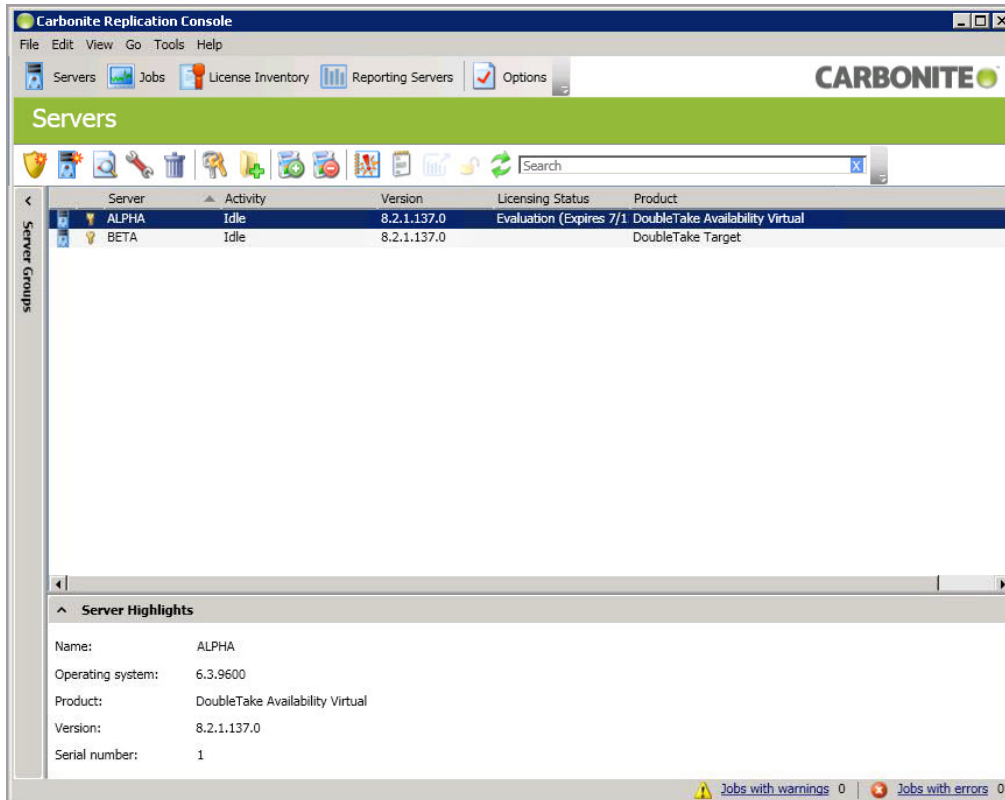
18. SQL Server 2005 or later may not initialize empty space when the database size increases due to the auto grow feature. Therefore, there is nothing for Carbonite Availability to replicate when this empty space is created. When the empty space is populated with data, the data is replicated to the target. 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 Availability will replicate the tag, unless it is a known driver. If it is a known driver, for example Microsoft SIS, Carbonite Availability 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 protection. You must load the reparse driver on the target after failover in order to access the data. Additionally, you cannot have reparse points in your data set if you are using same server protection because the server is functioning as both a source and target.
21. If you are using an archiving solution, do not archive any files after failover. Archiving files after failover could cause corruption.
22. If you are using the Microsoft Windows Update feature, keep in mind the following caveats.
 - a. Schedule your Windows Update outside the times when a mirroring operation (initial mirror, remirror, or a restoration mirror) is running. 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 Availability to only mirror changed files on reboot.
 - b. You must resolve any Windows Update incomplete operations or errors before failover or failback. (Check the windowsupdate.log file.) Also, do not failover or failback if the target is waiting on a Windows Update reboot. If failover occurs before the required Windows Update reboot, the target may not operate properly or it may not boot. You could also get into a situation where the reboot repeats indefinitely. One possible workaround for the reboot loop condition is to access a command prompt through the Windows Recovery Environment and delete the file \Windows\winsxs\pending.xml file. You may need to take ownership of the file to delete it. Contact technical support for assistance with this process or to evaluate other alternatives. Before you contact technical support, you should use the Microsoft System Update Readiness Tool as discussed in [Microsoft article 947821](#). This tool verifies and addresses many Windows Update problems.
23. If you are using Windows deduplication, keep in mind the following caveats.
 - a. Deduplication is only supported with files and folders and full server jobs. It is not supported with SQL jobs, provisioned virtual machine jobs, or image based jobs.
 - b. 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.
 - c. 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.
 - d. 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.

- e. After failover, the amount of disk space on the failed over server will be incorrect until you run the deduplication garbage collection which will synchronize the disk space statistics.
24. Replication is not case-sensitive. For example, if you rename the file Test.txt to test.txt, that change will not be replicated to the target. You will have to delete the file on the target and when it is remirrored, the new case of the file name will be used.
 25. 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 Availability 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 Availability 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 Availability 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 Availability 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 Availability 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 Availability servers.
 - **Default port for XML web services protocol**—Specifies the port that the console will use when sending and receiving data to Carbonite Availability 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 Availability version, you will need to use the legacy protocol port. This applies to Carbonite Availability 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 Availability 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 Availability 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 Availability where the installation files will be copied and run.
 - **Installation folder**—Specify the location where you want to install Carbonite Availability on each server. This field is not used if you are upgrading an existing version of Carbonite Availability. In that case, the existing installation folder will be used.
 - **Queue folder**—Specify the location where you want to store the Carbonite Availability 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 Availability 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 Availability will automatically begin the auto-disconnect process.
 - **Unlimited disk queue**—Carbonite Availability 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 Availability 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 Availability disk queuing. When the disk space limit is reached, Carbonite Availability 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 Availability 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 Availability 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 Availability 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 Availability source or target server which could be a physical server, virtual machine, or a cluster node
-  Carbonite Availability source or target server which is a Windows cluster
-  vCenter server
-  ESX server
-  Carbonite Availability 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 Availability 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 47.

Version

The Carbonite Availability 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 Availability 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 Availability on the server.

Product

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

Version

The product version information, if any

Serial Number

The serial number associated with the Carbonite Availability 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 Availability 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 43.

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 47.

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 49.

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 46. 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 a full server job with reverse protection enabled, you need to update the target image stored on the source if you change the credentials on the target server. See *Viewing full server job details* on page 183.

If you are using domain credentials for your Carbonite Availability 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 Availability servers using **Provide Credentials** in order to terminate the repeated pop-ups.

Manage Group Assignments

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

Install

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

Uninstall

Uninstalls Carbonite Availability 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 Availability 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 47.

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 49.

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 46. 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 a full server job with reverse protection enabled, you need to update the target image stored on the source if you change the credentials on the target server. See *Viewing full server job details* on page 183.

If you are using domain credentials for your Carbonite Availability 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 Availability 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 Availability on the selected server. This button opens the **Install** page where you can specify installation options.

Uninstall

Uninstalls Carbonite Availability 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 Availability 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 Availability log files; Carbonite Availability 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 protect 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 creating a files and folders job for identity failover (non-DNS), to best prepare for a potential failover, restore, and fallback scenario, you may want to use a source server that has been inserted into the console by a reserved IP address. In this scenario, your source must have at least two IP addresses, one for public communication and one for private. The private communication address will be the Carbonite Availability reserved address. If you insert your source in the console using the reserved IP address, then that reserved IP address can more easily be used after a failure to restore the data that changed during failover from the target back to the source. If your source server has already been inserted into the console by name, you can remove it and reinsert it by reserved IP address.

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 Availability 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 Availability 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 Availability.
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 Availability 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 Availability 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 Availability 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 Availability protocol being used to communicate with server. The protocol will be XML web services protocol (for servers running Carbonite Availability 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.

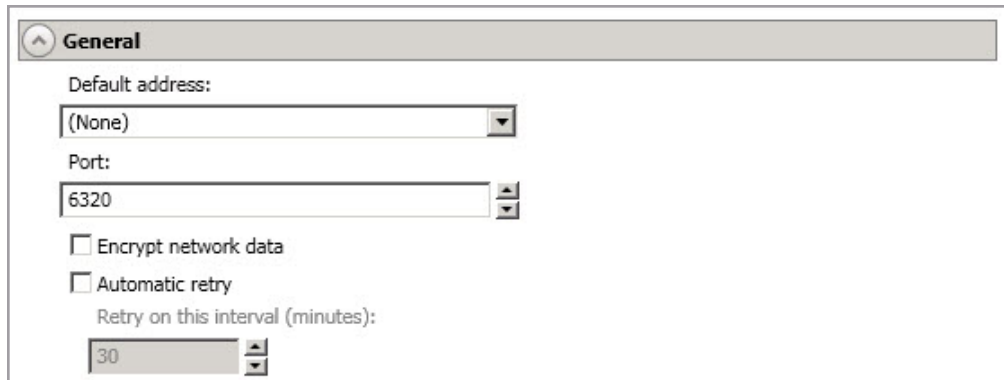
- *General server properties* on page 50—Identifies the server and configures encryption
- *Server licensing* on page 51—Views, adds, and removes license keys
- *Server setup properties* on page 54—Indicates how the server will act on startup and shutdown
- *Carbonite Availability queue* on page 57—Configures the Carbonite Availability queues
- *Source server properties* on page 61—Configures the source server
- *Target server properties* on page 63—Configures the target server
- *E-mail notification configuration* on page 65—Configures e-mail notification
- *Script credentials* on page 67—Specifies credentials to be used when executing custom scripts during mirroring or failover
- *Log file properties* on page 68—Configures log files



Server properties cannot be edited on a cluster.

General server properties

The general server properties identify the server and allow you to set encryption.



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 text input field containing "6320" with up and down arrow buttons on the right.
- Encrypt network data**
- Automatic retry**
- Retry on this interval (minutes):** A text input field containing "30" with up and down arrow buttons on the right.

- **Default address**—On a server with multiple NICs, you can specify which address Carbonite Availability 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 Availability 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 Availability 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.

The screenshot shows the 'Licensing' section of the Carbonite Replication Console. It includes a text input field for adding license keys with a 'Choose...' button, an 'Add' button, and a table of current license keys. Below the table are 'Remove' and 'Copy' buttons. There is also an 'Activation' section with a warning icon and instructions, a 'Server information' field with a 'Copy' button, and a link to 'Get an activation key from our website' and an 'Activate Online' button. At the bottom, there is an 'On Demand Licensing' section with fields for 'Service provider account number', 'Customer name', 'Product' (a dropdown menu), an 'Enable On Demand Proxy' checkbox, and a 'Proxy address' field. A 'Submit' button is at the very bottom.

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 Availability will not run.



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

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

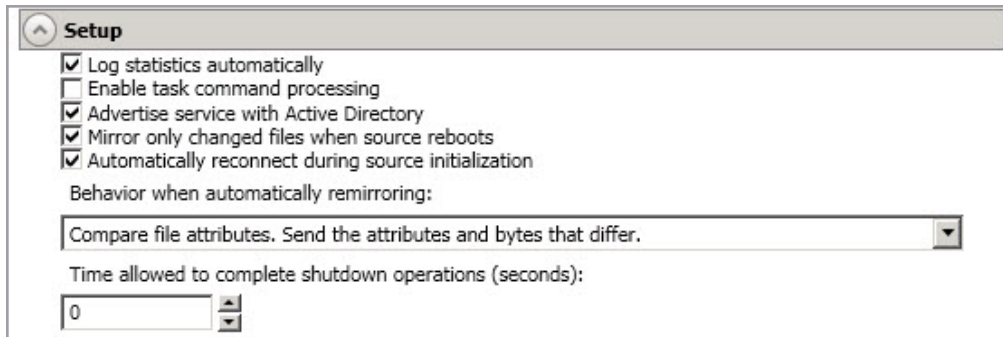
- **On Demand Licensing**—If you are a service provider participating in the On Demand licensing program, you can configure the subscription license for your target servers here. If you are not in this program, you can skip this section. For the latest and complete details on On Demand, see the help link in the On Demand web portal.
 1. Specify your **Service provider account number**. The account number is displayed in the upper right corner of the On Demand web portal.
 2. Specify the **Customer name**. Use the customer name configured on the Customers list in the On Demand web portal.
 3. Select the appropriate **Product** that corresponds with the Carbonite Availability product being used.
 4. If you are using a proxy server, select **Enable On Demand Proxy** and specify the **Proxy address** using the value `http://xxx.xxx.xxx.xxx:yyyy` where `xxx.xxx.xxx.xxx` is the IP

address of your proxy server and yyyy is the port number.

5. Click **Submit** to activate the subscription 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 Availability statistics logging will start automatically when Carbonite Availability is started.
- **Enable task command processing**—Task command processing is a Carbonite Availability 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 Availability 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 Availability will use the Carbonite Availability 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 Availability 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 Availability will automatically disconnect jobs when necessary. If this option is enabled, Carbonite Availability 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 Availability will automatically reconnect any jobs that were previously connected. Then, if configured, Carbonite Availability 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 Availability will automatically start disconnecting jobs. This is called auto-disconnect. The transaction logs and system memory are flushed allowing Carbonite Availability to begin processing anew. The auto-reconnect process ensures that any jobs that were auto-disconnected are automatically reconnected. Then, if configured, Carbonite Availability 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 Availability will auto-reconnect. If configured, Carbonite Availability 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 Availability 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 Availability 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 Availability source, not being able to communicate with the Carbonite Availability 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 Availability is able to successfully persist this information, when the Double-Take service on the target is restarted, Carbonite Availability will pick up where it left off, without requiring an auto-disconnect, auto-reconnect, or auto-remirror. If Carbonite Availability 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 Availability 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 Availability 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 Availability 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 Availability will perform the mirror when it is automatically remirroring.
-



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


- **Do not compare files. Send the entire file.**—Carbonite Availability 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.
 - **Compare file attributes. Send the entire file.**—Carbonite Availability will compare file attributes and will mirror those files that have different attributes, sending the entire file.
 - **Compare file attributes. Send the attributes and bytes that differ.**—Carbonite Availability will compare file attributes and will mirror only the attributes and bytes that are different.
 - **Compare file attributes and data. Send the attributes and bytes that differ.**—Carbonite Availability will compare file attributes and the file data and will mirror only the attributes and bytes that are different.
- **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 Availability 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 Availability. It does not control the service shutdown through a reboot or from the Service Control Manager.

Carbonite Availability queue

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

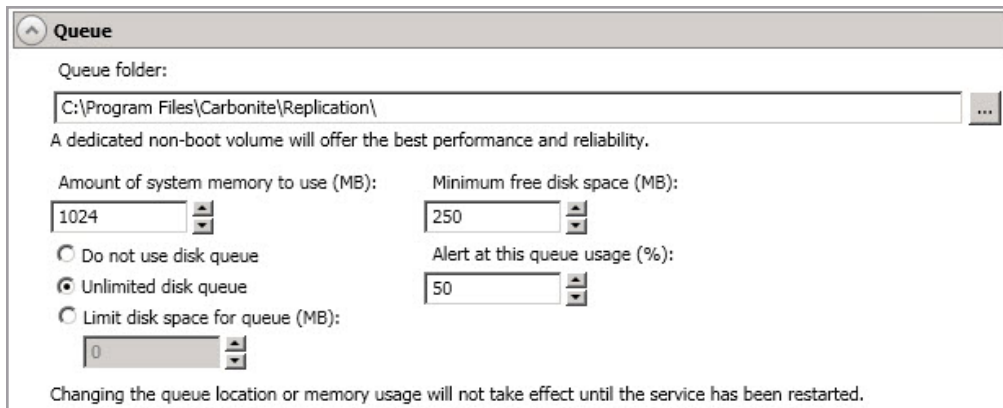
3. When system memory is full, the most recent changed data is added to the disk queue, as described in step 2. This means that system memory contains the oldest data. Therefore, when

data is transmitted to the target, Carbonite Availability 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 Availability 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 Availability 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 Availability pulls the data from system memory. Carbonite Availability 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 Availability server.



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

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

- Select an NTFS volume. Do not select a ReFS volume.
- Select a dedicated, non-boot volume.
- Do not select the same physical or logical volume as the data being replicated.
- Do not select the root of a 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.
- On a Windows 2012 or later 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 Availability 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 Availability queues.

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

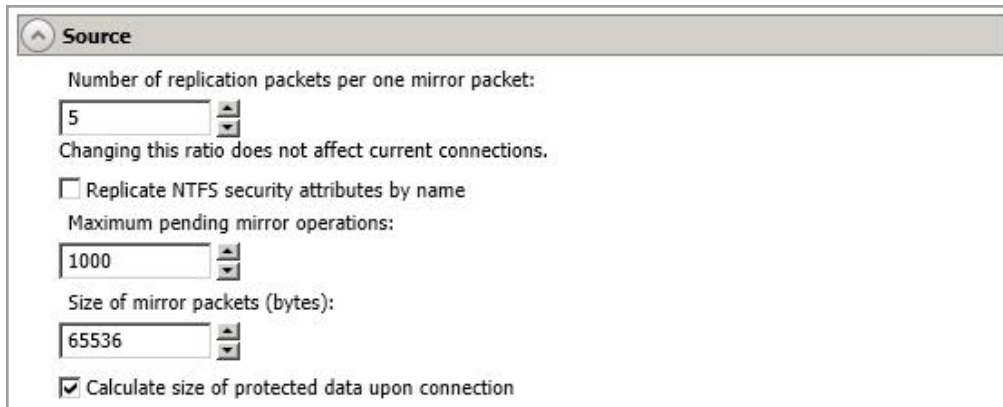


The **Limit disk space for queue** and **Minimum free disk space** settings work in conjunction with each other. For example, assume your queue is stored on a 10 GB disk with the **Limit disk space for queue** set to 10 GB and the **Minimum free disk space** set to 500 MB. If another program uses 5 GB, Carbonite Availability 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 Availability 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 Availability 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 Availability 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 Availability file and SID is replicated, a local user will not be able to access the file because the user's SID on the target machine is different from the SID that was replicated from the source machine.

If you enable this option, make sure that the same groups and users exist on the target as they do on the source. Additionally, you must enable this option on your target server before starting a restoration, because the target is acting like a source during a restoration.

Enabling this option may have an impact on the rate at which Carbonite Availability 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.

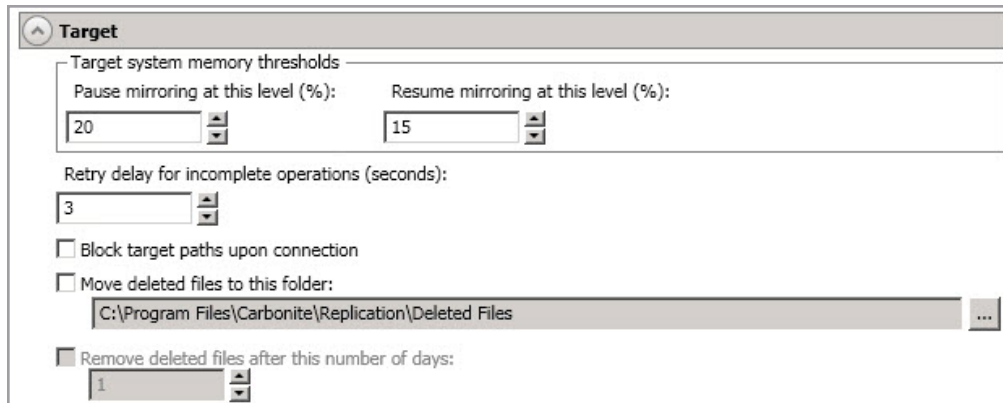
- **Maximum pending mirror operations**—This option is the maximum number of mirror operations that are queued on the source. The default setting is 1000. If, during mirroring, the mirror queued statistic regularly shows low numbers, for example, less than 50, this value can be increased to allow Carbonite Availability to queue more data for transfer.
- **Size of mirror packets**—This option determines the size of the mirror packets, in bytes, that Carbonite Availability 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 Availability 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 Availability processing. After failover, any target paths that are blocked will be unblocked automatically during the failover process so that users can modify data on the target after failover. During restoration, the paths are automatically blocked again. If you failover and failback without performing a restoration, the target paths will remain unblocked.



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

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

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



If you are moving deleted files on the target and you have orphan files configured for removal (which is the default setting for most job types), do not move the deleted files to a location inside the replica data on the target. The deleted files that are moved will then be deleted by the orphan file functionality.

- **Remove deleted files after this number of days**—If you are moving deleted files, you can specify a length of time, in days, to maintain the moved files. A moved file that is older than the specified number of days will be deleted. Carbonite Availability 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 Availability 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.

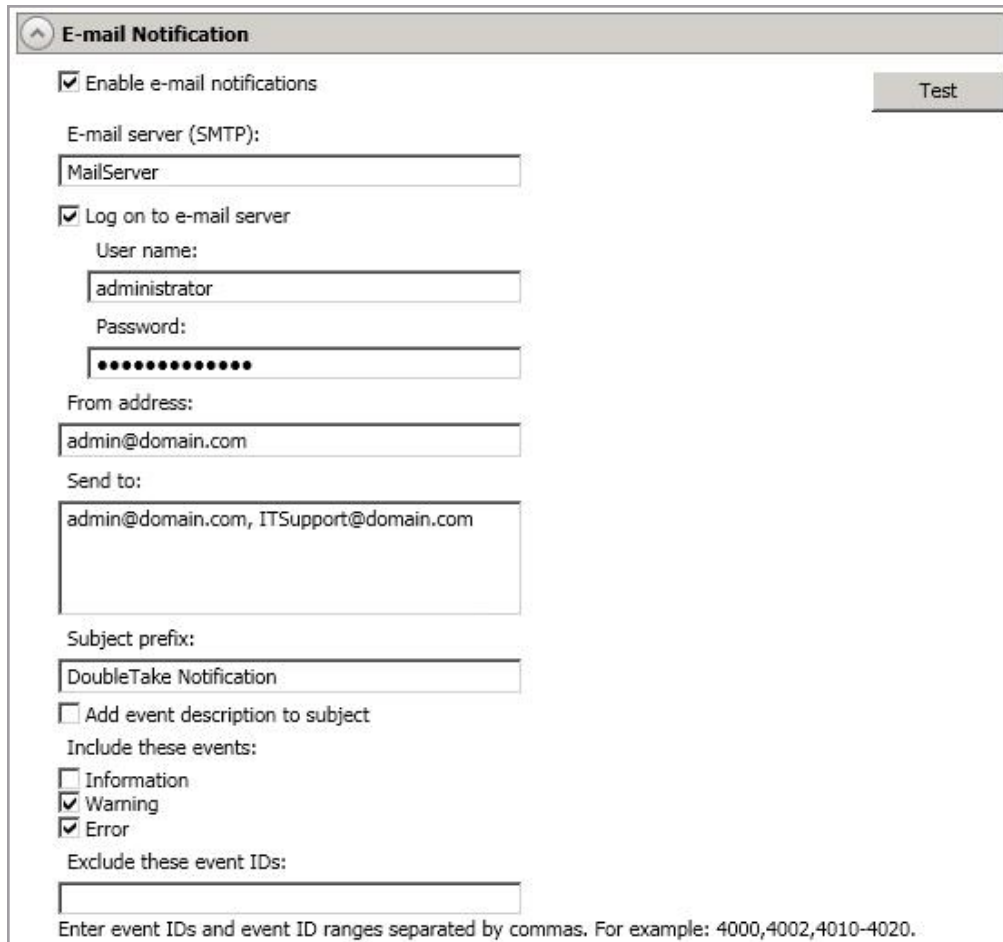
Alternate data streams that are deleted on the source will not be moved on the target.

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



E-mail Notification

Enable e-mail notifications Test

E-mail server (SMTP):
MailServer

Log on to e-mail server

User name:
administrator

Password:
●●●●●●●●

From address:
admin@domain.com

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

Subject prefix:
DoubleTake Notification

Add event description to subject

Include these events:

Information
 Warning
 Error

Exclude these event IDs:

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 Availability 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 Availability e-mail message. The address is limited to 256 characters.

- **Send to**—Specify the e-mail addresses that each Carbonite Availability 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 Availability e-mail message. This will help distinguish Carbonite Availability 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.

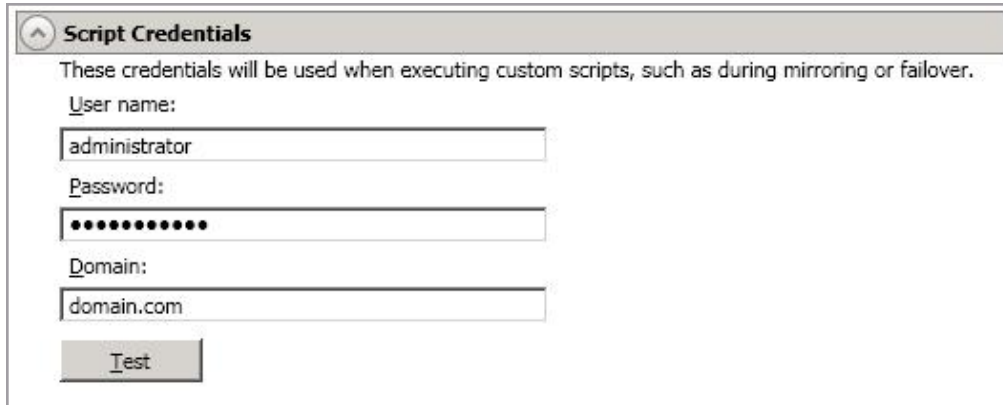
If an error occurs while sending an e-mail, a message will be generated. This message will not trigger another e-mail. Subsequent e-mail errors will not generate additional messages. When an e-mail is sent successfully, a message will then be generated. If another e-mail fails, one message will again 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 Availability e-mail messages are not blocked.

Script credentials

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



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 Availability 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 70.
 - **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 Availability 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: 1/24/2020 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 = 1/17/2020 8:21:36 PM Size = 1272 Mask = [0x20]
Target Attributes: Timestamp = 1/17/2020 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 = 1/17/2020 8:21:37 PM Size = 1272 Mask = [0x20]
Target Attributes: Timestamp = 1/17/2020 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 = 1/12/2020 3:28:20 PM Size = 17 Mask = [0x20]
Target Attributes: Timestamp = 1/20/2020 5:05:26 PM Size = 2 Mask = [0x20]
17 BYTES OUT OF SYNC
Completion Time: 1/24/2020 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.

Managing snapshots

Use the instructions below to manage the snapshots that Carbonite Availability has taken.

1. From the **Jobs** page, highlight the job and click **Manage Snapshots** in the toolbar.
2. You will see the list of snapshots, if any, associated with the job.
 - **Manual**—A user manually took this snapshot.
 - **Automatic**—Carbonite Availability automatically took this snapshot.
 - **Scheduled**—A periodic snapshot schedule triggered this snapshot.
 - **Deferred**—A periodic snapshot scheduled triggered this snapshot, although it did not occur at the specified interval because the job between the source and target was not in a good state.
 - **Test**—The test failover process took this snapshot.
 - **Coordinated**—A user took a coordinate snapshot.
 - **SQLClusterAutomatic**—The test failover process took this snapshot for a clustered SQL job.
3. Click **Take Snapshot** to create a new snapshot for the job.
4. If there is a snapshot that you no longer need, highlight it in the list and click **Delete**.
5. When you have completed your snapshot management, click **Close**.



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

Snapshot states

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

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

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

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

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

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

Chapter 5 Files and folders protection

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

- *Files and folders requirements* on page 81—Files and folders protection includes specific requirements for this type of protection.
- *Creating a files and folders job* on page 87—This section includes step-by-step instructions for creating a files and folders job.
- *Managing and controlling files and folders jobs* on page 114—You can view status information about your files and folders jobs and learn how to control these jobs.
- *Failing over files and folders jobs* on page 133—Use this section when a failover condition has been met or if you want to failover manually.
- *Failback and restoration for files and folders jobs* on page 134—Use this section to determine if you want to failback and then restore or if you want to restore then failback.



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

Files and folders requirements

Use these requirements for files and folders protection.

- **Operating system**—The following operating systems are supported for files and folders jobs.
 - Windows 2019 and Server Core 2019
 - 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 and 2019 support are for the primary operating system features available in Windows 2012. New operating system features specific to Windows 2016 and 2019, such as Nano Server, Windows Containers, and so on, are not supported.

Due to VMware limitations, Windows 2016 and 2019 are not supported on ESX 5.5.

DNS updates are not supported for Server Core servers.

- **File system**—Carbonite Availability supports the NTFS file system. On Windows 2016 and later, 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 21.
- **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 Availability 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 Availability queuing, logging, and so on.

- **Server name**—Carbonite Availability includes Unicode file system support, but your server name must still be in ASCII format. Additionally, all Carbonite Availability servers must have a unique server name.



If you need to rename a server that already has a Carbonite Availability license applied to it, you should deactivate that license before changing the server name. That includes rebuilding a server or changing the case (capitalization) of the server name (upper or lower case or any combination of case). If you have already rebuilt the server or changed



the server name or case, you will have to perform a host-transfer to continue using that license. See the *Carbonite Availability and Carbonite Migrate Installation, Licensing, and Activation* document for complete details.

- **Time**—The clock on your Carbonite Availability servers must be within a few minutes of each other, relative to UTC. Large time skews (more than five minutes) will cause Carbonite Availability 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 console 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 Availability 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 Availability.
 - 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 failover occurs the DHCP client will not start and an IP address cannot be assigned.
- **NAT support**—Carbonite Availability 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.
 - 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 Availability 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 Availability 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 Availability installation program will automatically attempt to configure ports 6320, 6325, and 6326 for Carbonite Availability. 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 367 for instructions on handling firewalls in your environment.

- **Windows Management Instrumentation (WMI)**—Carbonite Availability is dependent on the WMI service. If you do not use this service in your environment, contact technical support.
- **Carbonite Availability version**—You must have the same Carbonite Availability version on the source and target in order to perform a restoration.
- **Snapshots**—Support for Carbonite Availability snapshots for files and folders jobs is limited. You can take snapshots, however you cannot failover to a snapshot. The snapshots can be accessed and reverted on the target manually using VSS tools and utilities. Additionally, snapshots are not supported if your source and/or target is a cluster.

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

- **Snapshot location**—Snapshots are taken at the volume level and stored on the target. For example, if your job is protecting D:\data and E:\files, the snapshot will contain all of the data on both the D: and E: volumes. If your job is only protecting D:\data (E:\files exists but is not included in the job), the snapshot will only contain the D: volume. Make sure you have enough space on your target for snapshots.

- **Carbonite Availability installation location**—In order to enable Carbonite Availability snapshots, Carbonite Availability must be installed on the system drive. If Carbonite Availability is not installed on the system drive, snapshots will be disabled when enabling protection.
- **Server IP address**—If you have specified an IP address as the source server name, but that IP address is not the server's primary IP address, you will have issues with snapshot functionality. If you need to use snapshots, use the source's primary IP address or its name.
- **Snapshot limitations**—Sometimes taking a snapshot may not be possible. For example, there may not be enough disk space to create and store the snapshot, or maybe the target is too low on memory. If a snapshot fails, an Event message and a Carbonite Availability log message are both created and logged.

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

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

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

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

Volume Shadow Copy snapshots are associated with the volume they belong to. Since Carbonite Availability 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.

- **Clusters**—If you are using a cluster, make sure your cluster meets the following requirements.
 - **Best practices**—You should carefully review Microsoft documentation and resources for properly configuring your cluster before implementing Carbonite Availability on a cluster. The Microsoft TechNet articles [Failover Clusters](#) and [Installing and Upgrading on Cluster Nodes](#) are two resources you can start with. There are many other resources available on the [Microsoft TechNet web site](#).

- **Networking**—The following networking requirements apply to your cluster.
 - You must have TCP/IP connections between nodes.
 - Multiple networks are recommended to isolate public and private traffic.
 - The private network should be a unique subnet so that Carbonite Availability will not attempt to use an unreachable private network.
 - Your network can contain direct LAN connections or VLAN technology.
 - The source cluster IP must be accessible from the target.
- **Domain**—The cluster nodes must be members of the same domain.
- **DNS**—Forward and reverse lookups must be implemented on the primary DNS server for the cluster name and individual nodes.
- **Carbonite Availability disk queue**—Ensure that the disk queue is not on a Physical Disk resource.
- **Volumes**—The source and target should have identical drive mappings.
- **Owning nodes**—In a cluster configuration, if you add a possible owning node to the protected network name after a job has started, you must stop and restart the job. If you do not, the records for the new node will not be locked. This could cause problems with DNS records if the source cluster nodes are rebooted or the resources are otherwise cycled on the new owning node.
- **Licensing**—Each node in the cluster must have a valid Carbonite Availability license key.
- **Resource registration**—In some cases, the Carbonite Availability cluster resources may not be registered automatically when Carbonite Availability 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.
- **Supported configurations**—The following table identifies the supported configurations for a files and folders job.

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

Server Configuration	Description	Supported	Not Supported
One to many	You can protect a single source to multiple target servers. The source is the only server actively replicating data. This will create redundant copies of your source.	X	
Chained	You can protect 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. The first source and the middle server are the only servers actively replicating data.	X	
Single server	You can protect 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	Your servers can be in a standalone to cluster configuration, however this configuration does not support failover.	X	
Cluster to standalone	Your servers can be in a cluster to standalone configuration.	X	
Cluster to cluster	Your servers can be in a cluster to cluster configuration.	X	

If you are using a Cluster Shared Volume (CSV), you cannot protect any data, including a virtual machine, residing on the CSV. If you want to protect a CSV virtual machine, you must run Carbonite Availability 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 job

Use these instructions to create a files and folders job.



Note the following information concerning cluster scenarios.

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

1. If you are using a cluster to cluster configuration, follow these instructions to configure the target cluster before you create your Carbonite Availability files and folders job.
 - a. You may want to lower the TTL (time to live) value of your DNS records to avoid caching during this configuration.
 - b. Give the target cluster account full control over the source file server account in Active Directory and over the source file server DNS record.
-



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

- c. Take the source file server Client Access Point offline, including IP address, and then flush DNS on the target.
- d. Create the Client Access Point with the same name on the target and provide a unique IP address.
- e. While the target Client Access Point is online, create identical shares to the source file server.
- f. Take the target Client Access Point and IP address offline. Take the target file share resources offline.

- g. Bring the source Client Access Point online. DNS registration will change the IP address back to the source IP address.
 - h. If desired, reset your TTL value back to its original value.
 - i. If it is not already, DNS registration must be enabled for a source cluster name for DNS failover to function properly. You should not allow the source cluster name to come online if your source fails over.
 2. From the **Servers** page, right-click the server you want to protect and select **Protect**. You can also highlight a server, click **Create a New Job** in the toolbar, then select **Protect**.
 3. Choose the type of workload that you want to protect. Under **Server Workloads**, in the **Workload types** pane, select **Files and Folders**. In the **Workload items** pane, you will see the volumes and shares (if any) for your source, or if you are protecting a cluster, you will see your cluster groups. Select the volumes, shares, or cluster groups that you want to protect. You can select your files and folders in more detail in the **Replication Rules** section.
-



New shares within the original replication rules created after the job is created will be included on the target. However, new shares created outside of the original replication rules after the job is created will not be included on the target. For example, if C:\a is originally protected and a new share for C:\a\b is created, that share will be included on the target. However if C:\b is created, you must modify the replication rules to make sure the new share is included on the target.

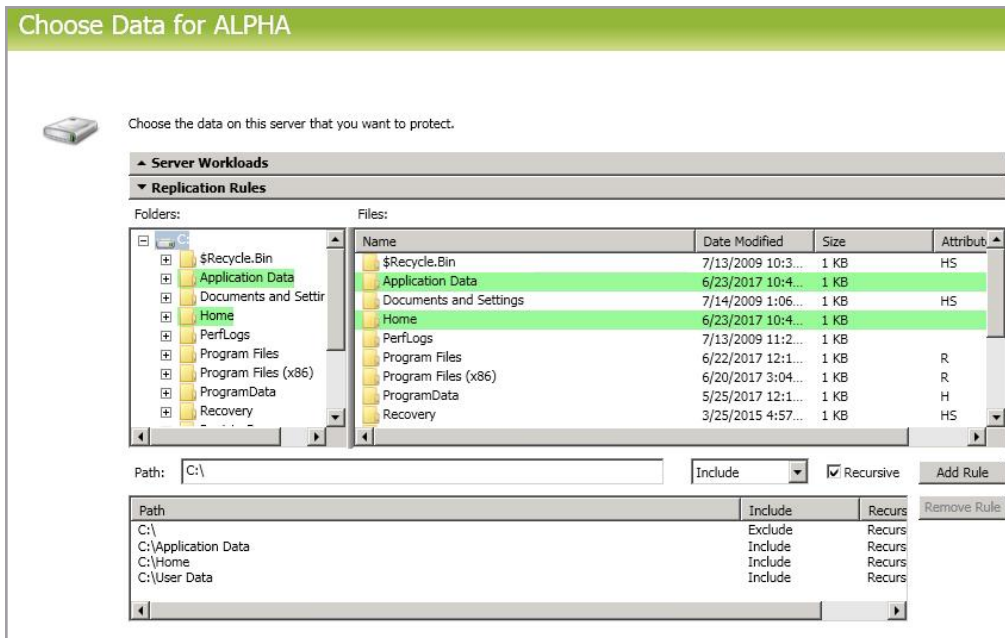
For clustered sources there are two options when selecting your cluster groups. One provides failover support (for data and server protection) and the other only DR support (for data only protection). You will not see the failover specific sections on the **Set Options** page later in the job creation workflow if you are only protecting data.

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.

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



If your source is a cluster and you are protecting data on a volume that is hosted by a node that does not currently own the cluster, you will not see the volume in the **Folders** list on the **Choose Data** page. You will have to manually enter your replication rules on this page.



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 Availability 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.

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



- **Current Servers**—This list contains the servers currently available in your console session. Servers that are not licensed for the workflow you have selected and those not applicable to the workload type you have selected will be filtered out of the list. Select your target server from the list. 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. Ideally, you should use a local account rather than a domain account because the domain account will fail to authenticate while failed over if the NetBIOS name and/or SPNs are failed over. If you want Carbonite Availability to update DNS during failover, the account must be a member of the Domain Admins group. If your security policies do not allow use of this group, see *DNS* on page 375 and use the instructions under the Carbonite Availability DFO utility to use a non-Domain Admins account.

7. 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.

8. You have many options available for your files and folders job. However, not all of the options will be applicable to clustered environments or to single server configurations. Configure those options that are applicable to your environment.

Go to each page identified below to see the options available for that section of the **Set Options** page. After you have configured your options, continue with the next step on page 112.



If your source is a cluster and you selected the DR only cluster groups (for data protection only), you will not see the failover specific sections on the **Set Options** page.

- *General* on page 91
- *Failover Monitor* on page 92
- *Failover Options* on page 94
- *Failover Identity* on page 97
- *Mirror, Verify & Orphaned Files* on page 100
- *Network Route* on page 103
- *Path Mapping* on page 104
- *Failover Services* on page 106
- *Snapshots* on page 107
- *Compression* on page 108
- *Bandwidth* on page 109
- *Scripts* on page 111

General

General

Job name:
alpha to beta

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

Failover Monitor

Failover Monitor

Total time to failure: 00:05:00

Consecutive failures: 20

Monitor on this interval: 00:00:10

Network monitoring

Monitor these addresses:

Source IP Address
<input checked="" type="checkbox"/> 172.31.206.200

Monitoring method: Network service

Failover trigger: All monitored IP addresses fail

Service monitoring

Services to monitor:

Add...
Remove

Attempt to restart this service after each failure

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

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

- **Consecutive failures**—Specify how many attempts the target will make to contact the source before the source is considered failed. For example, if you have this option set to 20, and your source fails to respond to the target 20 times in a row, this will be considered a failure.
- **Monitor on this interval**—Specify, in hours:minutes:seconds, how long to wait between attempts to contact the source to confirm it is online. This means that after a response (success or failure) is received from the source, Carbonite Availability will wait the specified interval time before contacting the source again. If you set the interval to 00:00:00, then a new check will be initiated immediately after the response is received.

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

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

- **Network monitoring**—With this option, the target will monitor the source using a network ping. Disable this option if you are using a standalone to cluster configuration, because failover is not supported.
 - **Monitor these addresses**—Select each **Source IP Address** that you want the target to monitor. If you are protecting a cluster, you are limited to the IP addresses in the cluster group that you are protecting. If you are using a NAT environment, do not select a private IP address on the source because the target cannot reach the source's private address, thus causing an immediate failure. Also for NAT environments, you will see an additional field for the **Replication Service port**. This gives you the ability to specify the port number to be used with the address, allowing the target to monitor the source through a router.
 - **Monitoring method**—At this time, **Network service** is the only option. Failover monitoring will test source availability by using an ICMP ping to confirm the route is active.
 - **Failover trigger**—If you are monitoring multiple IP addresses, specify when you want a failover condition to be triggered.
 - **One monitored IP address fails**—A failover condition will be triggered when any one of the monitored IP addresses fails. If each IP address is on a different subnet, you may want to trigger failover after one fails.
 - **All monitored IP addresses fail**—A failover condition will be triggered when all monitored IP addresses fail. If there are multiple, redundant paths to a server, losing one probably means an isolated network problem and you should wait for all IP addresses to fail.
- **Service monitoring**—This option is only available in standalone environments. It is not available in cluster environments. With this option, the target will monitor specific services on the source by confirming that they are running. Multiple services in the list will be checked in parallel. A failover condition is met when one of the monitored services fails the check. Click **Add** and select the service that you want to monitor. Repeat this step for additional services that you want to monitor. If you want to remove a service from the **Services to monitor** list, highlight it and click **Remove**.
 - **Attempt to restart this service after each failure**—When this option is enabled, if a service fails the monitor check, Carbonite Availability will attempt to restart it. During this restart period, Carbonite Availability will not check the other services, to avoid any false failures while the one service is attempting to be restarted. If the service cannot be restarted, Carbonite Availability will consider this a failure.

Failover Options



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

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

The screenshot shows the 'Failover Options' configuration window. It has a title bar with a collapse icon and the text 'Failover Options'. Below the title bar, there are several sections:

- Failover Options:** Contains two checked checkboxes: 'Wait for user to initiate failover' and 'Failover shares'.
- Active Directory:** Contains two unchecked checkboxes: 'Failover host name' and 'Failback host name'. Below these is a button labeled 'Active Directory Credentials...'. This section is enclosed in a rectangular box.
- Target scripts:** This section is divided into three sub-sections:
 - Pre-failover script:** Includes a text input field, a '...' button, and an 'Arguments:' label with an empty text input field. Below it is an unchecked checkbox labeled 'Delay failover until script completes'.
 - Post-failover script:** Includes a text input field, a '...' button, and an 'Arguments:' label with an empty text input field.
 - Pre-failback script:** Includes a text input field, a '...' button, and an 'Arguments:' label with an empty text input field. Below it is an unchecked checkbox labeled 'Delay failback until script completes'.
- Source scripts:** Includes a 'Post-failback script:' label, a text input field, a '...' button, and an 'Arguments:' label with an empty text input field.

- **Wait for user to initiate failover**—The failover process can wait for you to initiate it, allowing you to control when failover occurs. When a failure occurs, the job will wait in **Failover Condition Met** for you to manually initiate the failover process. Disable this option if you want failover to occur immediately when a failure occurs.
- **Failover shares**—Select this option to failover shares to the target. Only the shares that you selected on the **Choose Data** page will be protected and failed over.



Share failover only occurs for standard Windows file system shares. Other shares must be configured for failover through the failover scripts or created manually on the target. See *Macintosh shares* on page 385 or *NFS Shares* on page 386 for more information.

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

If your target is a standalone server, Windows share information is automatically updated on the target every 30 seconds. Since shares are not failed over if your target is a cluster, you will need to manually update shares on a cluster target. You will need to repeat the initial cluster configuration steps, although you can skip giving the target cluster account full control and creating the target Client Access Point, because those steps were already completed during the initial cluster configuration.

-
- **Failover host name**—If desired, you can failover the source server's host name. This will automatically remove the host SPN (Service Principal Name) from Active Directory on the source and add it to Active Directory on the target. If you are using Active Directory, enable this option or you may experience problems with failover. This option is not available for clustered jobs.
 - **Failback host name**—This option returns the host SPN on the source and target back to their original settings on failback. If you are using Active Directory, enable this option or you may experience problems with failback. This option is not available for clustered jobs.
 - **Active Directory Credentials**—If you are failing over and/or failing back the host name, you need to specify a user that has update privileges within Active Directory. Click **Active Directory Credentials** and identify a user and the associated password that has privileges to create and delete SPNs. The username must be in the format `fully_qualified_domain\user`, and the account password cannot be blank.
 - **Scripts**—You can customize failover and failback by running scripts on the source and target. Scripts may contain any valid Windows command, executable, or batch file. The scripts are processed using the same account running the Double-Take Management service, unless you have identified a specific account through the server's properties. See *Script credentials* on page 67. Examples of functions specified in scripts include stopping services on the target before failover because they may not be necessary while the target is standing in for the source, stopping services on the target that need to be restarted with the source's machine name and/or IP address, starting services or loading applications that are in an idle, standby mode waiting for failover to occur, notifying the administrator before and after failover or failback occurs, stopping services on the target after failback because they are no longer needed, stopping services on the target that need to be restarted with the target machine's original name and/or IP address, and so on. There are four types of failover and failback scripts that run on the target and one failback script that runs on the source.
 - **Pre-failover script**—This script runs on the target at the beginning of the failover process. Specify the full path and name of the script file.
 - **Post-failover script**—This script runs on the target at the end of the failover process. Specify the full path and name of the script file.
 - **Pre-failback script**—This script runs on the target at the beginning of the failback process. Specify the full path and name of the script file.
 - **Post-failback script**—This script runs on the target or source at the end of the failback process. Specify the full path and name of the script file.

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

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

Failover Identity

Failover Identity

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

Failover server name

Add these addresses to the selected target adapter after failover:

Source IP Address	Target Network Adapter
<input checked="" type="checkbox"/> 112.42.74.29	Local Area Connection
<input checked="" type="checkbox"/> 10.10.10.29	Local Area Connection 2

Retain target network configuration (Recommended for WAN configurations)

Failover server name (NetBIOS)

Update DNS server

DNS Options

Credentials for domain.com
User name: administrator

These DNS servers will be updated during failover:

112.42.48.9	<input type="button" value="Remove"/>
-------------	---------------------------------------

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

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

Update TTL (seconds):
300



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

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

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



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

Even though the **Apply source network configuration to the target** option is available for 2012 clusters, you must select **Retain target network configuration**. Because of changes to when resources in Windows 2012 clusters can come online related to DNS record locking, DNS updates are required for Windows 2012 cluster configurations, which is only an options when the retain target configuration option is selected.

- **Failover server name**—Select this option to failover the server name to the target. Carbonite Availability checks the hosts file and uses the first name there. If there is no hosts file, Carbonite Availability will use the first name in DNS. (Keep in mind, the first name in DNS may not always be the same each time the DNS server is rebooted.) Lastly, if there is no DNS server, Carbonite Availability will use the failover monitor name created by the Carbonite Replication Console. If you have selected to failover shares under the **Failover Options** sections, the server name will automatically be failed over so that the shares can be accessed after failover.
 - **Add these addresses to the selected target adapter after failover**—Select which IP addresses you want to failover and select the **Target Network Adapter** that will assume the IP address during failover.
-



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

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

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

- **Retain target network configuration**—If you select this option, the target will retain all of its original IP addresses. If your target is on a different subnet (typical of a WAN or NAT environment), you should select this option.
 - **Failover server name**—Select this option if you want to failover the NetBIOS name. This option will not be available for clustered targets.
 - **Update DNS server**—Specify if you want Carbonite Availability to update your DNS server on failover. If DNS updates are made, the DNS records will be locked during failover. Be sure and review the job requirements for updating DNS.
-



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

Because of changes to when resources in Windows 2012 clusters can come online related to DNS record locking, DNS updates are required for Windows 2012 cluster configurations.

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.

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

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



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

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

this case, you will need to specify failover scripts that update your DNS server during failover, or you can update the DNS server manually after failover.

Mirror, Verify & Orphaned Files

Mirror, 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 Availability 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 the mirror time 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 Availability 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 the mirror time 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 Availability 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 speed. 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 Availability 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 speed.

- **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 Availability 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 Availability 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 Availability 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.

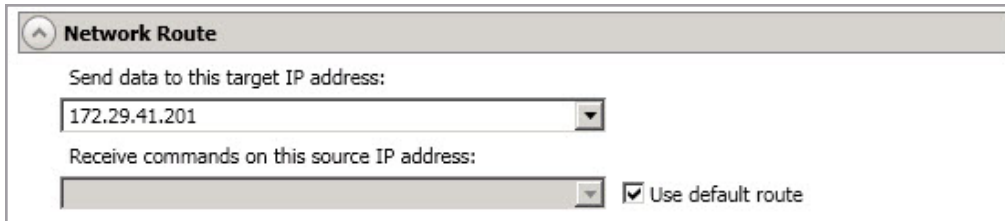
If delete orphaned files is enabled, carefully review any replication rules that use wildcard definitions. If you have specified wildcards to be excluded from protection, files matching those wildcards will also be excluded from orphaned file processing and will not be deleted from the target. However, 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 63 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 Availability 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 Availability log, and during difference mirrors, verifications, and restorations, all orphaned file processing messages are logged to the Carbonite Availability log. The orphaned file log is located in the **Logging folder** specified for the source. See *Log file properties* on page 68 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 this target IP address:
172.29.41.201

Receive commands on this source IP address:
 Use default route



If your target is a cluster, you will see a table presented where you can select the target route, rather than the single target route field as shown above. The field functions the same as described below.

- **Send data to this target IP address**—By default, Carbonite Availability will select an IP address on the target 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 Availability traffic. For example, you can separate regular network traffic and Carbonite Availability 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.

- **Receive commands on this source IP address**—By default, Carbonite Availability will select an IP address on the source to receive commands and requests for status from the target. This is communication from the Double-Take Management Service. If desired, specify an alternate route on the source that the commands and requests will be transmitted to. This allows you to select a different route for Carbonite Availability management communication. You can also manually enter a public IP address (which is the public IP address of the source server's router) if you are using a NAT environment.
- **Use default route**—Select this option to disable the drop-down list that allows you to select the route from the target server. When this option is enabled, the default route will automatically be used.

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



If you are protecting system state data (like your Program Files or Documents and Settings directory), you must select the **All to One** mapping or specify a customized location in order to avoid sharing violations. Keep in mind that this mapping will avoid sharing violations on the target, however during a restoration, you will get sharing violations on the source because the restoration mapping is one to one and your system state files will be in use on the source you are restoring to. In this case, restoration will never complete. If you will need to restore data and you must protect system state data, you should use a full server job.

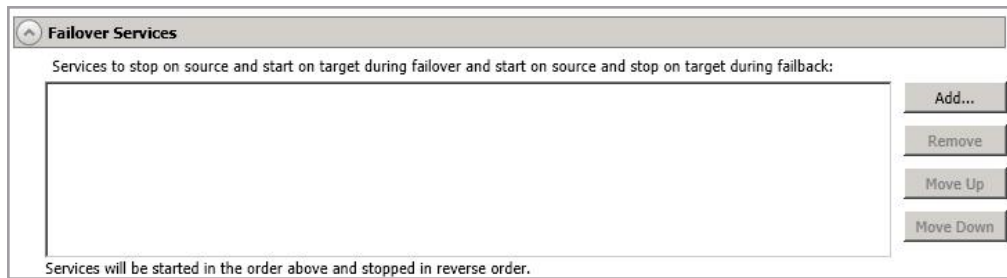
If you are protecting dynamic volumes or mount points, your location on the target must be able to accommodate the amount of data that may be stored on the source.

If you are protecting multiple mount points, your directory mapping must not create a cycle or loop. For example, if you have the C: volume mounted at D:\C and the D: volume mounted at C:\D, this is a circular configuration. If you establish a job for either C:\D or D:\C, there will be a circular configuration and Carbonite Availability 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 Availability processing. Any target paths that are blocked will be unblocked automatically during the failover process so that users can modify data after failover. During restoration, the paths are automatically blocked again. If you failover and failback without performing a restoration, the target paths will remain unblocked.

Failover Services



Failover Services

Services to stop on source and start on target during failover and start on source and stop on target during failback:

Services will be started in the order above and stopped in reverse order.

Add...
Remove
Move Up
Move Down



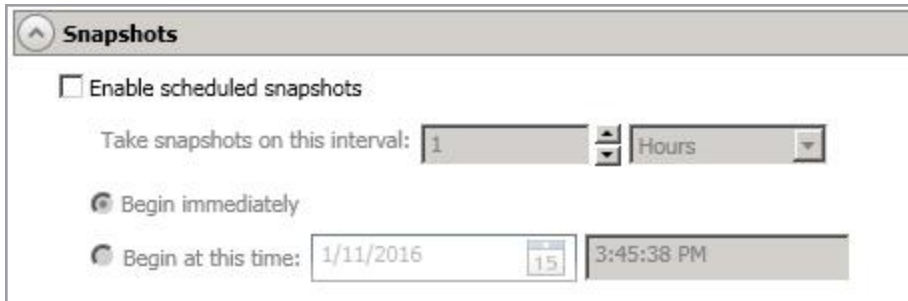
This section is not applicable to clustered environments.

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



You will have the option at failover time to override the stopping of options. If desired, you can select Leave source services running for live and snapshot failover.

Snapshots



This section is not applicable to clustered environments.

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

Turn on **Enable scheduled snapshots** if you want Carbonite Availability to take snapshots automatically at set intervals.


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



See *Managing snapshots* on page 76 for details on taking manual snapshots and deleting snapshots.

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

Compression



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

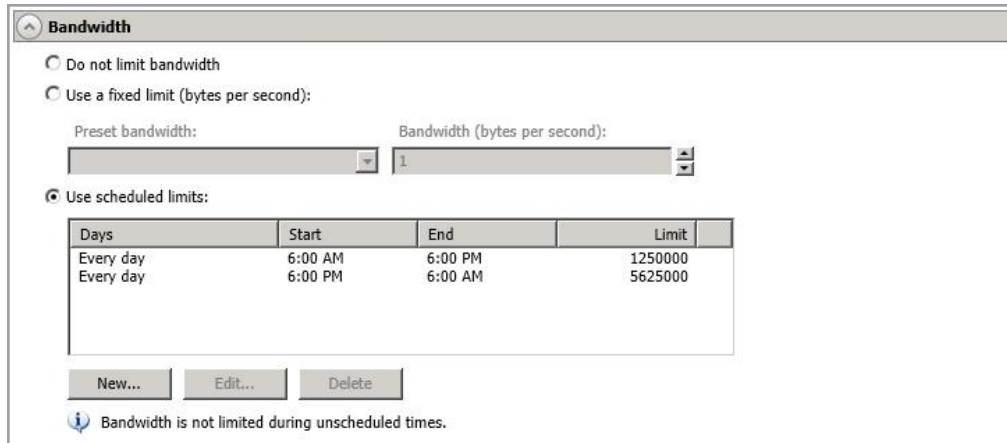
Keep in mind that the process of compressing data impacts processor usage on the source. If you notice an impact on performance while compression is enabled in your environment, either adjust to a lower level of compression, or leave compression disabled. Use the following guidelines to determine whether 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 Availability 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 Availability data transmissions. When a bandwidth limit is specified, Carbonite Availability never exceeds that allotted amount. The bandwidth not in use by Carbonite Availability is available for all other network traffic.



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

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

- **Do not limit bandwidth**—Carbonite Availability will transmit data using 100% bandwidth availability.
- **Use a fixed limit**—Carbonite Availability 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 Availability 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 configuration window titled "Scripts" with three sections: "Mirror Start", "Mirror Complete", and "Mirror Stop". Each section contains a "Script file" field, an "Arguments" field, and a "Test" button. Below each section are two checkboxes: "Allow script to interact with desktop" and "Delay until script completes".

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

If you are using Windows 2008 R2 with Service Pack 1, you can customize mirroring by running scripts on the target at pre-defined points in the mirroring process. Scripts are not functional for Windows 2012 and later.

Scripts may contain any valid Windows command, executable, or batch file. The scripts are processed using the same account running the Double-Take service, unless you have identified a specific account through the server's properties. See *Script credentials* on page 67. 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).

9. Click **Next** to continue.
10. Carbonite Availability 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 Availability correct the problem for you. For those errors that Carbonite Availability 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 Availability log files.

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



Jobs in a NAT environment may take longer to start.

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

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

Once a job is created, do not change the name of underlying hardware components used in the job. For example, volume names, network adapter names, or 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 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 33), 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 114
- *Detailed job information displayed in the bottom right pane* on page 117
- *Job controls* on page 119


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 Availability 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 a name or IP address of a standalone server, a cluster, or a node. Cluster jobs will be associated with the cluster name and standalone jobs will be associated with a standalone server or a cluster node.

Target Server

The name of the target. This could be a name or IP address of a standalone server, a cluster, or a node. Cluster jobs will be associated with the cluster name and standalone jobs will be associated with a standalone server or a cluster node.

Job Type

Each job type has a unique job type name. This job is a Files and Folders job. For a complete list of all job type names, press F1 to view the 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.
- **Restoring**—Data is being restored from the target to the source.
- **Unknown**—The console cannot determine the status.

Replication Status

- **Replicating**—Data is being replicated to the target.
- **Ready**—There is no data to replicate.
- **Pending**—Replication is pending.
- **Stopped**—Replication has been stopped.
- **Out of Memory**—Replication memory has been exhausted.
- **Failed**—The Double-Take service is not receiving replication operations from the Carbonite Availability 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.
- **Restore Required**—The data on the source and target do not match because of a failover condition. Restore the data from the target back to the source. If you want to discard the changes on the target, you can remirror to resynchronize the source and target.
- **Snapshot Reverted**—The data on the source and target do not match because a snapshot has been applied on the target. Restore the data from the target back to the source. If you want to discard the changes on the target, you can remirror to resynchronize the source and target.
- **Busy**—The source is low on memory causing a delay in getting the state of the data on the target.
- **Not Loaded**—Carbonite Availability 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 54.

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 54.

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. This is the current time where the console is running.

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 **Edit Job 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.

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 Availability queue was not exhausted during the time the job was paused. If the Carbonite Availability queue was exhausted during the time the job was paused, the job will restart mirroring and replication.

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

Pause

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

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

Stop

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

Take Snapshot

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

Snapshots are not applicable to clustered environments.

Manage Snapshots

Allows you to manage your snapshots by taking and deleting snapshots for the selected job. See *Managing snapshots* on page 76 for more information.

Snapshots are not applicable to clustered environments.

Failover or Cutover

Starts the failover process. See *Failing over files and folders jobs* on page 133 for the process and details of failing over a files and folders job.

Failback

Starts the failback process. See *Failback and restoration for files and folders jobs* on page 134 for the process and details of failing back a files and folders job.

Restore

Starts the restoration process. See *Failback and restoration for files and folders jobs* on page 134 for the process and details of restoring a files and folders job.

Reverse

Reverses protection. Reverse protection does not apply to files and folders jobs.

Undo Failover or Cutover

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

View Job Log

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

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 Availability 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 Availability 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 Availability 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 the mirror time 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 Availability 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 the mirror time 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 Availability 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 speed. 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 Availability 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 speed.
- **Calculate size of protected data before mirroring**—Specify if you want Carbonite Availability 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 Availability will transmit data using 100% bandwidth availability.
 - **Fixed bandwidth limit**—Carbonite Availability 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 Availability 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 Availability processing, not the entire server.

While the target is paused, the Carbonite Availability 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 Availability 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.

- **Refresh Status**—Refreshes the job status immediately.

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 33.

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 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 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.
- **Restore Required**—The data on the source and target do not match because of a failover condition. Restore the data from the target back to the source. If you want to discard the changes on the target, you can remirror to resynchronize the source and target.
- **Snapshot Reverted**—The data on the source and target do not match because a snapshot has been applied on the target. Restore the data from the target back to the source. If you want to discard the changes on the target, you can remirror to resynchronize the source and target.
- **Busy**—The source is low on memory causing a delay in getting the state of the data on the target.
- **Not Loaded**—Carbonite Availability 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 Availability 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.
- **Restoring**—Data is being restored from the target to the source.
- **Unknown**—The console cannot determine the status.

Mirror percent complete

The percentage of the mirror that has been completed

Mirror remaining

The total number of mirror bytes that are remaining to be sent from the source to the target. This value may be zero if you have enabled **Mirror only changed files when source reboots** on the *Server setup properties* on page 54.

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 54.

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 Availability 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 job

Over time, you may want to confirm that any changes in your network or environment have not impacted your Carbonite Availability 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 Availability 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 Availability correct the problem for you. For those errors that Carbonite Availability 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 job

Use these instructions to edit a files and folders 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 Availability monitor security access.)
2. You will see the same options available for your files and folders 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 files and folders job* on page 87 for details on each job option.



Changing some options may require Carbonite Availability 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 Availability 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 Availability orphan file detection cannot remove the data for you. Therefore, you have to remove it manually.

-
4. Click **Next** to continue.
 5. Carbonite Availability 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 Availability correct the problem for you. For those errors that Carbonite Availability 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 Availability log files.

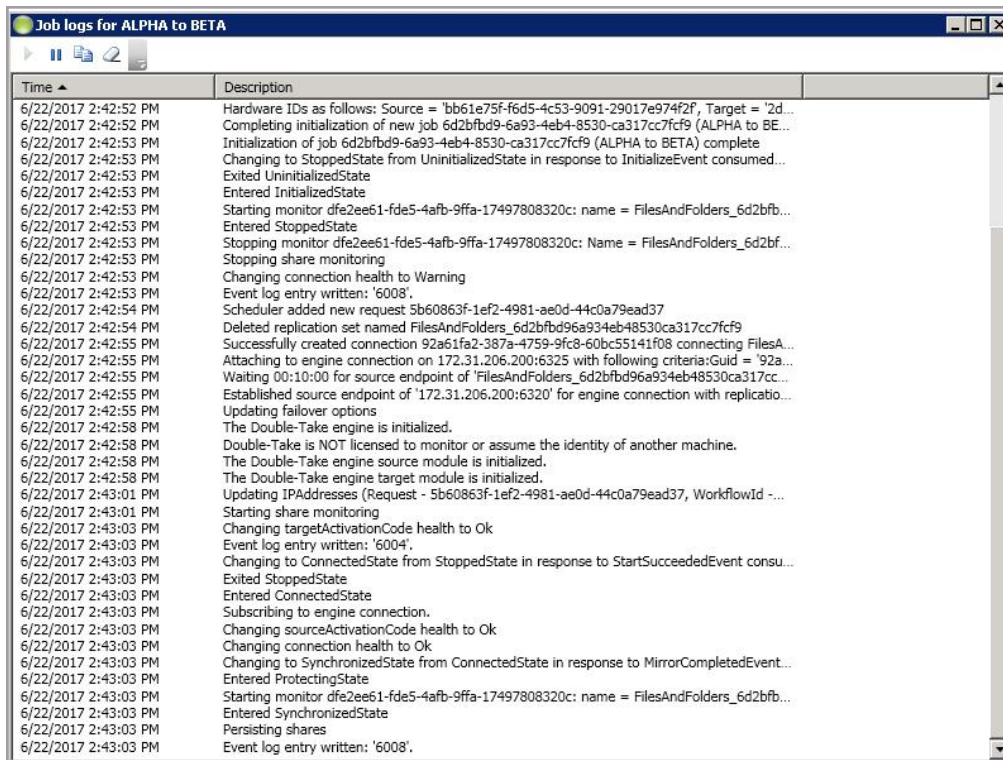
6. Once your servers have passed validation and you are ready to update your job, click **Finish**.

Viewing a files and folders job log

You can view a job log file through the 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.



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



Time	Description
6/22/2017 2:42:52 PM	Hardware IDs as follows: Source = 'bb1e75f-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_6d2bf...
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.

Failing over files and folders jobs

When a failover condition has been met, failover will be triggered automatically if you disabled the wait for user option during your failover configuration. If the wait for user before failover option is enabled, you will be notified in the console when a failover condition has been met. At that time, you will need to trigger it manually from the console when you are ready.



If you have paused your target, failover will not start if configured for automatic failover, and it cannot be initiated if configured for manual intervention. You must resume the target before failover will automatically start or before you can manually start it.

If you are using a files and folders job in a standalone to cluster configuration, the **Failover, Cutover, or Recover** button will be enabled once a mirror is complete. Do not failover as unexpected results may occur. Failover is not supported for files and folders jobs in a standalone to cluster configuration.

1. On the **Jobs** page, highlight the job that you want to failover and click **Failover or Cutover** in the toolbar.
2. Select the type of failover to perform.
 - **Failover to live data**—Select this option to initiate a full, live failover using the current data on the target. The target will stand in for the source by assuming the network identity of the failed source. User and application requests destined for the source server or its IP addresses are routed to the target.
 - **Perform test failover**—This option is not available for files and folders jobs.
 - **Failover to a snapshot**—This option is not applicable to files and folders jobs.
3. Select how you want to handle the data in the target queue.
 - **Apply data in target queues before failover or cutover**—All of the data in the target queue will be applied before failover begins. The advantage to this option is that all of the data that the target has received will be applied before failover begins. The disadvantage to this option is depending on the amount of data in queue, the amount of time to apply all of the data could be lengthy.
 - **Discard data in the target queues and failover or cutover immediately**—All of the data in the target queue will be discarded and failover will begin immediately. The advantage to this option is that failover will occur immediately. The disadvantage is that any data in the target queue will be lost.
4. The last option will vary depending on if your source is clustered.
 - **Leave source services running**—If your source is standalone, select this option if you want the services specified in **Failover options** in the job options to remain online instead of being shutdown. This is for live or snapshot failover. It does not apply to test failover.
 - **Leave source cluster resources online**—If your source is clustered, select this option if you want the cluster resources to remain online instead of being taken offline. This is for live or snapshot failover. It does not apply to test failover.
5. When you are ready to begin failover, click **Failover**.

Failback and restoration for files and folders jobs

Failover occurred because the target was monitoring the source for a failure, and when a failure occurred, the target stood in for the source. User and application requests that were directed to the failed source are routed to the target.

While the users are accessing their data on the target, you can repair the issue(s) on the source. Before users can access the source again, you will need to restore the data from the target back to the source and perform failback. Failback is the process where the target releases the source identity it assumed during failover. Once failback is complete, user and application requests are no longer routed to the target, but back to the source.

Ideally, you want to restore your data from the target back to the source before you failback. This allows users who are currently accessing their data on the target because of failover to continue accessing their data. Restoration before failback reduces user downtime. Another method, which may be easier in some environments that have strict IP addressing policies, allows you to failback first and then restore the data from the target to the source. A possible disadvantage to this process is that users may experience longer downtime, depending on the amount of data to be restored, because they will be unable to access their data during both the restoration and the failback.

- *Restoring then failing back files and folders jobs* on page 135
- *Failing back then restoring files and folders jobs* on page 137




If you want to failback only, without performing a restoration, follow the instructions for failing back then restoring and you will be able to skip the restoration process. Keep in mind that if you skip the restoration process, any data changes that were made on the target during failover will be lost.

Restoring then failing back files and folders jobs

Restoring before failing back allows your users to continue accessing their data on the failed over target, which is standing in for the source, while you perform the restoration process. The key to this process is to keep the users off of the source, but allow the source and target to communicate to perform the restoration.

1. Resolve the problem(s) on the source that caused it to fail. Make sure in resolving the problems, that you do not bring the source on the network at this time because the target currently has the source's identity because of failover.
2. Disable all of the IP addresses on the source that you failed over to the target.
3. If you failed over all of your source IP addresses, change an existing IP address on the source to a new, unique IP address that the target can access. If you inserted your source server into the console using a reserved IP address or a public NAT address when you created the job, and you did not failover that IP address, you can skip this step.
4. Configure your new, unique IP address that you created or the reserved IP address that you are using so that it does not automatically register with DNS. This option is on the Advanced TCP/IP Settings dialog box on the DNS tab.
5. Bring the source onto the network using the IP address that the target can access. You can disregard any identity conflict errors.
6. Stop any applications that may be running on the source. Files must be closed on the source so that updated files from the target will overwrite the files on the source.
7. If you had to create a new, unique IP address on the source, you will have to remove the original source in the console and add it back in using the new, unique IP address. Use a local account, not a domain account, that is a member of the Double-Take Admin and Administrators groups. Complete this step on the **Servers** page. If you inserted your source server into the console using a private IP address or a public NAT address when you created the job, you can skip this step.
8. On the **Jobs** page, highlight the job and click **Restore**.
9. Confirm **Restore data** is selected, then highlight your source server in the server list. If your server has a public NAT address, you can disable the default communication port and specify another port number to use, allowing the servers to communicate through a router.



10. Click **Continue** to start the restoration.



During the restoration, only the data is restored back to the source. Shares are not created on the source during the restoration. Shares that were created on the target during failover will need to be created manually on the source, if they do not already exist.

11. When the restoration is complete, highlight the job and click **Failback**.
12. In the dialog box, highlight the job that you want to failback and click **Failback**.



If you do not see your job after failback, remove and re-add your target to the console. The job may not be visible depending on where you are running the console from.

13. After failback is complete and the job is stopped, enable or add the IP addresses on the source that you disabled earlier. Make sure that you keep any new addresses that you created because the job is still using that address.
14. If you restored to a new source and are going to enable protection again, edit the job to reconfigure your failover settings.
15. Click **Start** to restart protection.

Failing back then restoring files and folders jobs


Failback before restoration can be a simpler process, but it may require additional downtime. The amount of downtime will depend on the amount of data to be restored. Users must be kept off of the source and target during this entire process.

1. Remove the source from the network and fix the issue that caused your source server to fail. Make sure in resolving the problems that you do not bring the source on the network at this time because the target currently has the source's identity because of failover.
2. Schedule a time for the failback and restoration process. Select a time that will have minimal disruption on your users.
3. When you are ready to begin the failback process, power on the source, if it is not on already. Make sure you do not connect it to the network. You must prohibit user access to both the source and target.
4. On the **Jobs** page, highlight the job and click **Failback**.
5. Highlight the job that you want to failback and click **Failback**.
6. Once failback is complete, connect the source to the network. Make sure that end users continue to be prohibited from accessing both the source and target because the updated data from the target needs to be restored back to the source.



Depending on where you are running the console from, you may need to add the target back to the console after failback in order to see your job.

7. Stop any applications that may be running on the source. Files must be closed on the source so that updated files from the target will overwrite the files on the source.
8. On the **Jobs** page, highlight the job and click **Restore**.



9. If you want to skip the restoration, select **Do not restore data**, and click **Continue**. Keep in mind that if you skip the restoration process, any data changes that were made on the target during failover will be lost. If you want to restore the changed data from the target back to the source, select **Restore data**, highlight your source server in the server list. If your server has a public NAT address, you can disable the default communication port and specify another port number to use, allowing the servers to communicate through a router.
10. Click **Continue** to start the restoration.



During the restoration, only the data is restored back to the source. Shares are not created on the source during the restoration. Shares that were created on the target during failover will need to be created manually on the source, if they do not already exist.

11. When the restoration job is complete, the job will automatically be stopped. At this point, you can allow users to access the source again.
12. If you restored to a new source and are going to enable protection again, edit the job to reconfigure your failover settings.
13. Click **Start** to restart protection.

Chapter 6 Full server protection

Create a full server job when you want to protect the entire source, including the server's system state. You can also use it to protect an application server. This type of job is the most flexible, allowing you to go from physical to physical, physical to virtual, virtual to virtual, and virtual to physical.

- *Full server requirements* on page 140—Full server protection includes specific requirements for this type of protection.
- *Creating a full server job* on page 149—This section includes step-by-step instructions for creating a full server job.
- *Managing and controlling full server jobs* on page 173—You can view status information about your full server jobs and learn how to control these jobs.
- *Failing over full server jobs* on page 192—Use this section when a failover condition has been met or if you want to failover manually.
- *Reversing full server jobs* on page 196—Use this section to reverse protection. The source (what was your original target hardware) is now sending data to the target (what was your original source hardware).

Full server requirements

Use these requirements for full server protection. Keep in mind that a target server may meet these requirements but may not be suitable to stand-in for a source in the event of a source failure. See *Target compatibility* on page 146 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 2019 and Server Core 2019
 - 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 and 2019 support are for the primary operating system features available in Windows 2012. New operating system features specific to Windows 2016 and 2019, such as Nano Server, Windows Containers, and so on, are not supported.

Due to VMware limitations, Windows 2016 and 2019 are 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 protection 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 failing 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.
 - **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 failover since there are thousands of extra files that need to be examined. Ideally, your target should be as clean and simple a configuration as possible.



During failover, Carbonite Availability 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 Availability supports the NTFS file system. On Windows 2016 and later, 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 21.
 - **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 Availability program files.
 - **Source storage**—If you will be enabling reverse protection, the source must have enough space to store, process, and apply the target's system state data. See the disk space requirements in *Target compatibility* on page 146 for details on the space requirements for various operating systems.
 - **Server name**—Carbonite Availability includes Unicode file system support, but your server name must still be in ASCII format. Additionally, all Carbonite Availability servers must have a unique server name.
-



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

- **Time**—The clock on your Carbonite Availability servers must be within a few minutes of each other, relative to UTC. Large time skews (more than five minutes) will cause Carbonite Availability 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 console 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 Availability 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 Availability.
 - If your source or target server only supports DHCP, for example Windows Azure, keep in

mind the following caveats.

- You will be unable to reverse.
 - 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 failover occurs the DHCP client will not start and an IP address cannot be assigned.
- **Network adapters**—Microsoft NIC teaming for Windows 2012 is supported, however no third party NIC teaming solutions are supported.
 - **NAT support**—Carbonite Availability 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 Availability 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 Availability 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 Availability installation program will automatically attempt to configure ports 6320, 6325, and 6326 for Carbonite Availability. 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 367 for instructions on handling firewalls in your environment.

- **Windows Management Instrumentation (WMI)**—Carbonite Availability is dependent on the WMI service. If you do not use this service in your environment, contact technical support.
- **Snapshots**—You can take and failover to Carbonite Availability snapshots using a full server job.

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

- **Snapshot location**—Snapshots are taken at the volume level and stored on the target. For example, if your job is protecting D:\data and E:\files, the snapshot will contain all of the data on both the D: and E: volumes. If your job is only protecting D:\data (E:\files exists but is not included in the job), the snapshot will only contain the D: volume. Make sure you have enough space on your target for snapshots.
- **Carbonite Availability installation location**—In order to enable Carbonite Availability snapshots, Carbonite Availability must be installed on the system drive. If Carbonite Availability is not installed on the system drive, snapshots will be disabled when enabling protection.
- **Server IP address**—If you have specified an IP address as the source server name, but that IP address is not the server's primary IP address, you will have issues with snapshot functionality. If you need to use snapshots, use the source's primary IP address or its name.
- **Snapshot limitations**—Sometimes taking a snapshot may not be possible. For example, there may not be enough disk space to create and store the snapshot, or maybe the target is too low on memory. If a snapshot fails, an Event message and a Carbonite Availability log message are both created and logged.

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

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

You must also keep in mind that if you are using extended functionality provided by Volume Shadow Copy, you need to be aware of the impacts that functionality may have on

Carbonite Availability. For example, if you change the location where the shadow copies are stored and an error occurs, it may appear to be a Carbonite Availability error when it is in fact a Volume Shadow Copy error. Be sure and review any events created by the VolSnap driver and check your Volume Shadow Copy documentation for details.

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

Volume Shadow Copy snapshots are associated with the volume they belong to. Since Carbonite Availability 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.

- **Supported configurations**—The following table identifies the supported configurations for a full server job.

Server Configuration	Description	Supported	Not Supported
One to one active/standby	You can protect 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 protect 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 protect many source servers to one target server.		X
One to many	You can protect a single source to multiple target servers. The source is the only server actively replicating data. This will create redundant copies of your source. You will only be able to configure reverse protection for the first job. Subsequent jobs from that source will have reverse protection disabled.	X	
Chained	You cannot protect 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

Server Configuration	Description	Supported	Not Supported
Single server	You cannot protect a single source to itself.		X
Standalone to standalone	Your servers can be in a standalone to standalone configuration.	X	
Standalone to cluster	Your servers cannot be in a standalone to cluster configuration.		X
Cluster to standalone	Your servers cannot be in a cluster to standalone configuration.		X
Cluster to cluster	Your servers cannot be in a cluster to cluster configuration.		X

If you are using a Cluster Shared Volume (CSV), you cannot protect any data, including a virtual machine, residing on the CSV. If you want to protect a CSV virtual machine, you must run Carbonite Availability 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 failover since there are thousands of extra files that need to be examined. Ideally, your target should be as clean and simple a configuration as possible.



During failover, Carbonite Availability 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 failing 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 failover. Also because Windows Azure only supports DHCP, you will be unable to reverse a full server job.
- **Server role**—The target cannot be a domain controller. Ideally, the target should not host any functionality (file server, application server, and so on) because the functionality will be removed when failover occurs.

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

Additionally, if your source is a domain controller, you will not be able to reverse protection.

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

- **Network adapters**—You must map at least one NIC from the source to one NIC on the target. If you have NICs on the source that are not being used, it is best to disable them. If the source has more NICs than the target, some of the source NICs will not be mapped to the target. Therefore, the IP addresses associated with those NICs will not be available after failover. If there are more NICs on the target than the source, the additional NICs will still be available after failover and will retain their pre-failover network settings.
- **File system format**—The source and the target must have the same NTFS 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 protecting on the source, the target must have a matching volume. For example, if you are protecting drives C: and D: on the source, the target cannot have drives D: and E:. In this case, the target must also have drives C: and D:. Additional target volumes are preserved and available after failover with all data still accessible. You will be unable to reverse protection if the target has more drives than the source.
- **System path**—The source and the target must have the same system path. The system path includes the location of the Windows files, Program Files, and Documents and Settings.
- **Carbonite Availability version**—If you will be using the reverse feature with your full server job, your source and target must be running the same Carbonite Availability version.
- **Disk space**—The target must have enough space to store the data from the source. This amount of disk space will depend on the applications and data files you are protecting. The more data you are protecting, the more disk space you will need. The target must also have enough space to store, process, and apply the source's system state data. If you will be enabling reverse protection, the source must have enough space to store, process, and apply the target's system state data. In either case, the size of the system state will depend on the operating system and architecture.

A copy of the source system state data will be staged on the target boot volume in a folder called Staging-SSM. For reverse protection, the target system state will be staged on the source boot volume also 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
- **Windows 2016**—at least 25 GB
- **Windows 2019**—at least 23 GB

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

Creating a full server job

Use these instructions to create a full server job.

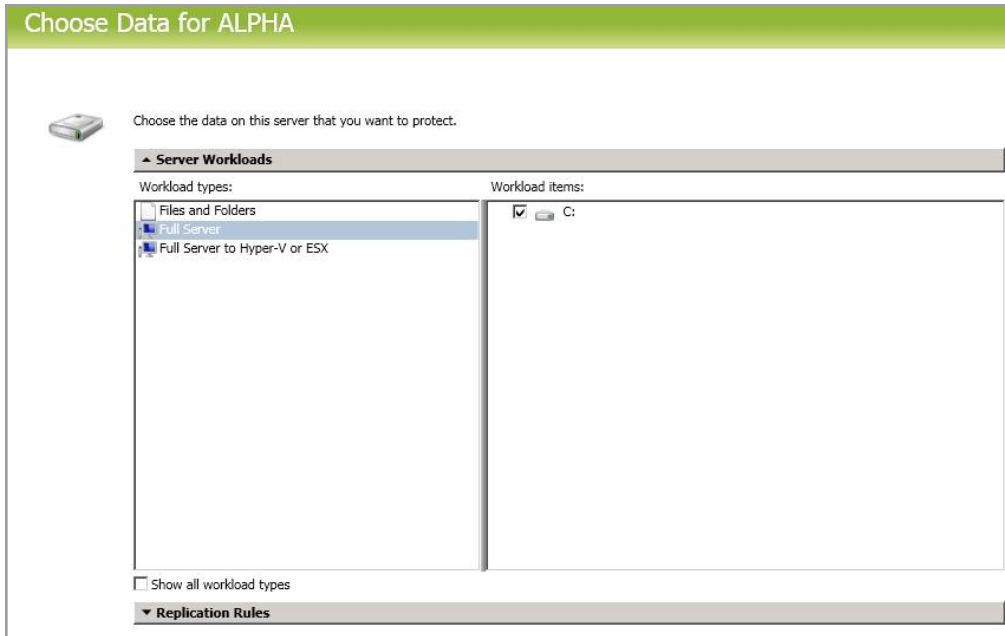


If you are creating a full server job in order to revert back to your original configuration after failing over a full server to ESX or full server to Hyper-V, you will need to perform a few additional tasks before creating the full server job. Contact technical support if you need assistance with these steps.

1. On either the source or target, stop the Double-Take and Double-Take Management Service services.
 2. Remove the GUID value from HKEY_LOCAL_MACHINE\SOFTWARE\NSI Software\Double-Take\CurrentVersion\Communication\Uniqueld. Do not delete the Uniqueld key. Only delete the GUI value within the key.
 3. Restart the the Double-Take and Double-Take Management Service services.
 4. Remove and then add your servers back into the Carbonite Replication Console.
 5. Install a different license on the original source and complete a host transfer if necessary.
-

1. Review these best practices before you create your job.
 - **NIC configuration**—If you are planning to failover the IP address of the source, use a separate NIC and separate network for a Carbonite Availability reserved IP address that will not be failed over. If you are unable to do that and just one NIC is used for both production and reserved IP addresses, disable DNS registration on the NIC. If you are not going to failover the IP address of the source, an additional NIC and address is not necessary. In this case, Carbonite Availability will block the DNS record for that address while it is failed over.
 - **Separate console machine**—Ideally, you should use a separate machine (not the source or target) to run the console, configure protection, to failover, and to reverse. The separate machine must be able to communicate with the source and target using their reserved IP addresses.
 - **Carbonite Replication Console**—Insert your source and target servers into the console using the reserved IP addresses and a local computer account that is a member of the Double-Take Admin and Administrators groups.
2. From the **Servers** page, right-click the server you want to protect and select **Protect**. You can also highlight a server, click **Create a New Job** in the toolbar, then select **Protect**.
3. Choose the type of workload that you want to protect. Under **Server Workloads**, in the **Workload types** pane, select **Full Server**. In the **Workload items** pane, select the volumes on the source that you want to protect.

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



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

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

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


If you need to remove a rule, highlight it in the list at the bottom and click **Remove Rule**. Be careful when removing rules. Carbonite Availability 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.

5. Click **Next** to continue.
6. Choose your target server. This is the server that will store the replica data from the source, and in the event of a failover, it will become your 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 the source is a domain controller, Carbonite Availability will add the security groups to the users OU, therefore permissions must be located there. If your source is the only domain controller in your network, the account must also be a local account in the local administrators group on the target. If you want Carbonite Availability to update DNS during failover, the account must be a member of the Domain Admins group. If your security policies do not allow use of this group, see *DNS* on page 375 and use the instructions under the Carbonite Availability DFO utility to use a non-Domain Admins account.

If you are protecting a Hyper-V source and you select an existing Hyper-V server to use as the target, the Hyper-V Integration Services on the target must be version 2008 SP2 or greater. Without this version, the target may not start after failover.

7. 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.

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

Go to each page identified below to see the options available for that section of the **Set Options** page. After you have configured your options, continue with the next step on page 172.

- *General* on page 153
- *Failover Monitor* on page 154
- *Test Failover* on page 155
- *Failover Options* on page 157
- *Failover Identity* on page 158
- *Reverse Protection and Routing* on page 160
- *Network Adapter Options* on page 162
- *Mirror, Verify & Orphaned Files* on page 162
- *Staging Folder Options* on page 165
- *Target Services* on page 166
- *Snapshots* on page 167
- *Compression* on page 168
- *Bandwidth* on page 169
- *Scripts* on page 171

General



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

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

Failover Monitor

Failover Monitor

Total time to failure: 00:05:00

Consecutive failures: 20

Monitor on this interval: 00:00:10

Network monitoring

Monitor these addresses:

	Source IP Address
<input checked="" type="checkbox"/>	172.31.206.201

Monitoring method: Network service

Failover trigger: All monitored IP addresses fail

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

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

- **Consecutive failures**—Specify how many attempts the target will make to contact the source before the source is considered failed. For example, if you have this option set to 20, and your source fails to respond to the target 20 times in a row, this will be considered a failure.
- **Monitor on this interval**—Specify, in hours:minutes:seconds, how long to wait between attempts to contact the source to confirm it is online. This means that after a response (success or failure) is received from the source, Carbonite Availability will wait the specified interval time before contacting the source again. If you set the interval to 00:00:00, then a new check will be initiated immediately after the response is received.

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

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

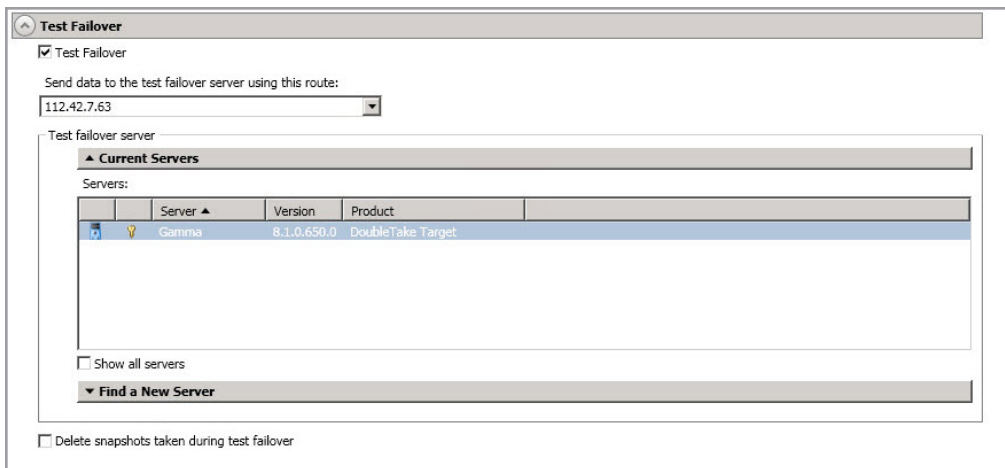
- **Network monitoring**—With this option, the target will monitor the source using a network ping.
 - **Monitor these addresses**—Select each **Source IP Address** that you want the target to monitor. If you are using a NAT environment, do not select a private IP address on the source because the target cannot reach the source's private address, thus causing an immediate failure. Also for NAT environments, you will see an additional field for the **Replication Service port**. This gives you the ability to specify the port number to be used with the address, allowing the target to monitor the source through a router.
 - **Monitoring method**—At this time, **Network service** is the only option. Failover monitoring will test source availability by using an ICMP ping to confirm the route is active.
 - **Failover trigger**—If you are monitoring multiple IP addresses, specify when you want a failover condition to be triggered.
 - **One monitored IP address fails**—A failover condition will be triggered when any one of the monitored IP addresses fails. If each IP address is on a different subnet, you may want to trigger failover after one fails.
 - **All monitored IP addresses fail**—A failover condition will be triggered when all monitored IP addresses fail. If there are multiple, redundant paths to a server, losing one probably means an isolated network problem and you should wait for all IP addresses to fail.

Test Failover

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

- The network adapters on the test failover machine must have the same friendly name as the network adapters you have mapped for live failover in the **Network Adapter Options** section. For example, if you have Local Area Network mapped to Local Area Network, the test failover machine must have an adapter named Local Area Network. If you have Local Area Network mapped to Network 2, and the test failover machine does not have an adapter named Network 2, test failover will fail.
- The source, target, and protection job will remain online and uninterrupted during the test.
- During the test, any scheduled snapshots for the protection job will be deferred until after the test server is online.
- The test will be performed using the test failover settings configured during job creation.
- The test failover will take a snapshot of the current data on the target and mirror the data from the snapshot to the test failover machine using the same mirroring options as the protection job.
- Once the mirror is complete, the test failover machine is rebooted automatically to finalize the test failover process.
- The test failover machine will maintain its own identity during the test.
- The test failover machine will maintain its own networking which keeps it isolated from the rest of the network in order to avoid network conflicts and redirecting clients. Applications or functionality that relies on the source networking may not be fully testable with the test machine networking.

- When using the test failover machine, log in with local credentials to avoid trust relationship issues.
- When you are finished with your test, undo it.
- When you undo a test failover, the snapshot 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, including your test failover snapshot. (Performing a live failover or failing over to a snapshot will automatically undo any test in progress.)



- **Test Failover**—Enable this option to be able to perform test failover.
- **Send data to the test failover server using this route**—Select or enter a route to use on the test failover server for mirroring the data from the snapshot to the test failover server.
- **Test failover server**—Select the server you want to use for the test failover.
 - **Current Servers**—This list contains the servers currently available in your console session. Servers that are not applicable to test failover will be filtered out of the list. Select your test failover 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. 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. After you have identified or located the server, click **Add Server**. If there are any issues connecting to that server, you will see an error in yellow at the top of the page. If there are no issues, you can continue.
- **Delete snapshots taken during test failover**—Select this option if you want to delete the snapshots taken for the test failover process. If you disable this option, the snapshots will not be deleted when you perform undo failover.

Failover Options

Failover Options

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

Target scripts

Pre-failover script: Arguments:

Delay failover until script completes

Post-failover script: Arguments:

- **Wait for user to initiate failover**—The failover process can wait for you to initiate it, allowing you to control when failover occurs. When a failure occurs, the job will wait in **Failover Condition Met** for you to manually initiate the failover process. Disable this option if you want failover to occur immediately when a failure occurs.
- **Change target ports to match source during failover**—This option allows you to specify if you want the Carbonite Availability ports on the target to be updated to match the source during failover. This option is useful in NAT environments where the public port does not match the private port.
- **Scripts**—You can customize failover by running scripts on the target or the recovered source. Scripts may contain any valid Windows command, executable, or batch file. The scripts are processed using the same account running the Double-Take Management service, unless you have identified a specific account through the server's properties. See *Script credentials* on page 67. Examples of functions specified in scripts include stopping services on the target before failover because they may not be necessary while the target is standing in for the source, stopping services on the target that need to be restarted with the source's machine name and/or IP address, starting services or loading applications that are in an idle, standby mode waiting for failover to occur, notifying the administrator before and after failover occurs, and so on. There are two types of failover scripts.
 - **Pre-failover script**—This script runs on the target at the beginning of the failover process. Specify the full path and name of the script file.
 - **Post-failover script**—This script runs on the recovered source at the end of the failover process. Specify the full path and name of the script file.
 - **Arguments**—Specify a comma-separated list of valid arguments required to execute the script.
 - **Delay until script completes**—Enable this option if you want to delay the failover process until the associated script has completed. If you select this option, make sure your script handles errors, otherwise the failover process may never complete if the process is waiting on a script that cannot complete.

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

Failover Identity

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**

These DNS servers will be updated during failover:

112.42.48.9	<input type="button" value="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 failover to the target. If your target is on the same subnet as the source (typical of a LAN environment), you should select this option. Do not select this option if you are using a NAT environment that has a different subnet on the other side of the router.



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

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

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



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

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

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.

Reverse Protection and Routing

Reverse Protection and Routing

Send data to this target IP address:
172.29.41.201

Receive commands on this source IP address:
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:
172.29.41.200

Select a reserved IP address on the target:
172.29.41.201

- **Send data to this target IP address**—By default, Carbonite Availability will select an IP address on the target 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 Availability traffic. For example, you can separate regular network traffic and Carbonite Availability 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 and for communication after a reverse. The **Replication Service port** is used for data transmission.
- **Receive commands on this source IP address**—By default, Carbonite Availability will select an IP address on the source to receive commands and requests for status from the target. This is communication from the Double-Take Management Service. If desired, specify an alternate route on the source that the commands and requests will be transmitted to. This allows you to select a different route for Carbonite Availability management communication. You can also manually enter a public IP address (which is the public IP address of the source server's router) if you are using a NAT environment.
- **Use default route**—Select this option to disable the drop-down list that allows you to select the route from the target server. When this option is enabled, the default route will automatically be used.
- **Enable reverse protection**—After failover, your target server is lost. Reverse protection allows you to store a copy of the target's system state on the source server, so that the target server will not be lost. The reverse process will bring the target identity back on the source hardware and establish protection. After the reverse, the source (running on the original target hardware) will be protected to the target (running on the original source hardware).

If you do not use reverse protection, after a failover, your target server will be lost. In order to continue protecting your data, you will have to manually rebuild your original source and restart protection, which can be a long and complicated process. Also, if you disable reverse, you will lose the activated target license after failover. See the *Carbonite*

Availability and Carbonite Migrate Installation, Licensing, and Activation document for details on how you can save and deactivate the target license.

You may want to consider having two IP addresses on each server. This will allow you to monitor and failover one (or more) IP addresses, while still leaving an IP address that does not get failed over. This IP address that is not failed over is called a reserved IP address and can be used for the reverse process. The reserved IP address remains with the server hardware. Ideally, the reserved IP address should not be used for production communications. The reserved IP address can be on the same or a different subnet from your production IP addresses, however if the subnet is different, it should be on a different network adapter. The reserved IP addresses will also be used to route Carbonite Availability data.

You do not have to have a second IP address on each server. It is acceptable to use the production IP address for reverse protection, as long as you are selecting the option to retain the target configuration.

- **Select a reserved IP address on the source**—Specify an IP address on the source which will be used to permanently identify the source server. The IP address you specify will not be failed over to the target in the event of a failure. This allows you to reverse protection back to the source after a failover.
- **Send data to source after reverse using this route**—This field will only be displayed if the console recognizes that your source address is a public NAT address. In that case, you can specify the route and disable the default communication port and specify another port number to use for data transmission.
- **Select a reserved IP address on the target**—Specify an IP address on the target which will be used to permanently identify the target server. The IP address you specify will not be lost during failover. This allows you to reverse protection back to the source after a failover. In a non-NAT environment, this address will override the target route above and be used to route the data to the target server.



When reverse protection is enabled, your source server must have space to store, process, and apply the target's system state data. See *Disk space* under *Full server requirements* on page 140 for estimates on the amount of space you may need.

When the job is first started and reverse protection is enabled, an image of the target's system state is mirrored to the source server. This mirror may cause a performance impact on your source server. This impact is only temporary, and system performance will return to normal when the reverse protection mirror is complete.

To maintain system performance on the source, the target's system state is not continuously replicated to the source. You can manually update the image of the target's system state by viewing the job details and clicking **Update** under **Target Server Image**. See *Viewing full server job details* on page 183.

Network Adapter Options

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. Also, if you have enabled reverse protection, make sure that your NICs with your reserved IP addresses are mapped to each other.

Mirror, Verify & Orphaned Files

Mirror 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 Availability 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 the mirror time 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 Availability 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 speed.

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 Availability 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 speed.
 - **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 Availability 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 Availability 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 Availability 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 63 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 Availability 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 Availability log, and during difference mirrors, verifications, and restorations, all orphaned file processing messages are logged to the Carbonite Availability log. The orphaned file log is located in the **Logging folder** specified for the source. See *Log file properties* on page 68 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.

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 Availability will be unable to write to open application files. In this case, you will want to mirror those application files to a staging location instead of their actual location. Generally, this will only apply to applications that are not installed in the Windows Program Files directory. In this case, click **Add** and specify the folder that you want staged. Any staged folders will be applied to their actual installation location during failover.



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

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

Target Services



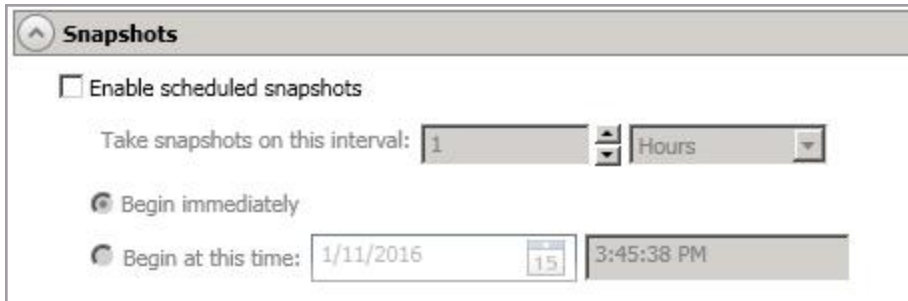
- **Services to leave running on the target server during protection**—Carbonite Availability controls which services are running and stopped on the target during protection. You can specify which services you want to keep running by clicking **Add** and selecting a service from the list. If you want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



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

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

Snapshots



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

Turn on **Enable scheduled snapshots** if you want Carbonite Availability to take snapshots automatically at set intervals.


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



See *Managing snapshots* on page 76 for details on taking manual snapshots and deleting snapshots.

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

Compression



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

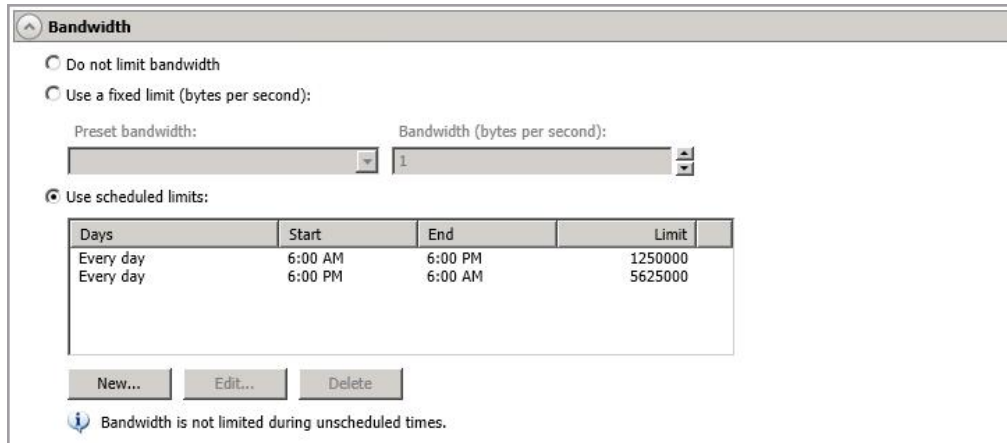
Keep in mind that the process of compressing data impacts processor usage on the source. If you notice an impact on performance while compression is enabled in your environment, either adjust to a lower level of compression, or leave compression disabled. Use the following guidelines to determine whether 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 Availability 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 Availability data transmissions. When a bandwidth limit is specified, Carbonite Availability never exceeds that allotted amount. The bandwidth not in use by Carbonite Availability 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 Availability will transmit data using 100% bandwidth availability.
- **Use a fixed limit**—Carbonite Availability 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 Availability 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 configuration window titled "Scripts" with three sections: "Mirror Start", "Mirror Complete", and "Mirror Stop". Each section contains a "Script file" field, an "Arguments" field, and a "Test" button. Below each section are two checkboxes: "Allow script to interact with desktop" and "Delay until script completes".

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

If you are using Windows 2008 R2 with Service Pack 1, you can customize mirroring by running scripts on the target at pre-defined points in the mirroring process. Scripts are not functional for Windows 2012 and later.

Scripts may contain any valid Windows command, executable, or batch file. The scripts are processed using the same account running the Double-Take service, unless you have identified a specific account through the server's properties. See *Script credentials* on page 67. 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).

9. Click **Next** to continue.
10. Carbonite Availability 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 Availability correct the problem for you. For those errors that Carbonite Availability 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 Availability log files.

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



Jobs in a NAT environment may take longer to start.

Once a job is created, do not change the name of underlying hardware components used in the job. For example, volume names, network adapter names, or 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 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 33), 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 173
- *Detailed job information displayed in the bottom right pane* on page 176
- *Job controls* on page 178

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 Availability 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. In parenthesis will be the reserved IP address that you assigned to this server.

Target Server

The name of the target. This could be the name or IP address of your target. In parenthesis will be the reserved IP address that you assigned to this server.

Job Type

Each job type has a unique job type name. This job is a Full Server 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 Availability 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.
- **Snapshot Reverted**—The data on the source and target do not match because a snapshot has been applied on the target. Restore the data from the target back to the source. If you want to discard the changes on the target, you can remirror to resynchronize the source and target.
- **Busy**—The source is low on memory causing a delay in getting the state of the data on the target.
- **Not Loaded**—Carbonite Availability 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 54.

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 54.

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. This is the current time where the console is running.

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 **Edit Job 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.

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 Availability queue was not exhausted during the time the job was paused. If the Carbonite Availability queue was exhausted during the time the job was paused, the job will restart mirroring and replication.

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

Pause

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

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

Stop

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

Take Snapshot

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

Manage Snapshots

Allows you to manage your snapshots by taking and deleting snapshots for the selected job. See *Managing snapshots* on page 76 for more information.

Snapshots are not applicable to clustered environments.

Failover or Cutover

Starts the failover process. See *Failing over full server jobs* on page 192 for the process and details of failing over a full server job.

Failback

Starts the failback process. Failback does not apply to full server jobs.

Restore

Starts the restoration process. Restore does not apply to full server jobs.

Reverse

Reverses protection. The original source hardware will be reversed to the target identity and the job will start mirroring in the reverse direction with the job name and log file names changing accordingly. After the mirror is complete, the job will continue running in the opposite direction. See *Reversing full server jobs* on page 196 for the process and details of reversing a full server job.

The backup job, when a full server protection job is configured for reverse, may become orphaned after the initial backup job has been completed, manually updated, and then the target server is restarted. In this specific case, you will need to contact technical support for a workaround to use the backup job during a reverse.

Undo Failover or Cutover

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

View Job Log

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

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 Availability 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 Availability 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 Availability 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 the mirror time 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 Availability 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 speed. 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 Availability 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 speed.
- **Calculate size of protected data before mirroring**—Specify if you want Carbonite Availability 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 Availability will transmit data using 100% bandwidth availability.
 - **Fixed bandwidth limit**—Carbonite Availability 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 Availability 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 Availability processing, not the entire server.

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



If you have paused your target, failover will not start if configured for automatic failover, and it cannot be initiated if configured for manual intervention. You must resume the target before failover will automatically start or before you can manually start it.

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

- **Refresh Status**—Refreshes the job status immediately.

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 33.

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 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 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.
- **Snapshot Reverted**—The data on the source and target do not match because a snapshot has been applied on the target. Restore the data from the target back to the source. If you want to discard the changes on the target, you can remirror to resynchronize the source and target.
- **Busy**—The source is low on memory causing a delay in getting the state of the data on the target.
- **Not Loaded**—Carbonite Availability 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 Availability 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 54.

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 54.

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 Availability 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

Target Server Image

When a full server job is created with reverse protection enabled, an image of the target's system state is stored on the source server. This image allows you to reverse your source and target after a failover. To improve performance, the target's system state is not continuously replicated to the source. You should manually update the image of the target's system state by clicking **Update** if there is a change on the target. For example, if the credentials on the target server are updated, you should update the target server image that is on the source. This reverse protection mirror may cause a performance impact on your source server. This impact is only temporary, and system performance will return to normal when the reverse protection mirror is complete.

If you have reverse enabled, are updating your target image, and the Double-Take service on the target restarts (either manually or automatically, for example the target server restarts), you should restart your target image update after the Double-Take service is back online. This will correct any incorrect status displayed in the console and ensure the target image is complete.

Validating a full server job

Over time, you may want to confirm that any changes in your network or environment have not impacted your Carbonite Availability 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 Availability 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 Availability correct the problem for you. For those errors that Carbonite Availability 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 job

Use these instructions to edit a full server 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 Availability monitor security access.)
2. You will see the same options for your full server job as when you created the job, but you will not be able to edit all of them. If desired, edit those options that are configurable for an existing job. See *Creating a full server job* on page 149 for details on each job option.



Changing some options may require Carbonite Availability 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 Availability 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 Availability orphan file detection cannot remove the data for you. Therefore, you have to remove it manually.

-
4. Click **Next** to continue.
 5. Carbonite Availability 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 Availability correct the problem for you. For those errors that Carbonite Availability 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 Availability log files.

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

Viewing a full server job log

You can view a job log file through the 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.



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

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_6d2bf...
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.

Failing over full server jobs

When a failover condition has been met, failover will be triggered automatically if you disabled the wait for user option during your failover configuration. If the wait for user before failover option is enabled, you will be notified in the console when a failover condition has been met. At that time, you will need to trigger it manually from the console when you are ready.



If you have paused your target, failover will not start if configured for automatic failover, and it cannot be initiated if configured for manual intervention. You must resume the target before failover will automatically start or before you can manually start it.

1. On the **Jobs** page, highlight the job that you want to failover and click **Failover or Cutover** in the toolbar.
2. Select the type of failover to perform.
 - **Failover to live data**—Select this option to initiate a full, live failover using the current data on the target. The source is automatically shut down if it is still running. Then the target will stand in for the source by rebooting and applying the source identity, including its system state, on the target. After the reboot, the target becomes the source, and the target no longer exists.
 - **Perform test failover**—Select this option to perform a test failover.
 - The network adapters on the test failover machine must have the same friendly name as the network adapters you have mapped for live failover in the **Network Adapter Options** section. For example, if you have Local Area Network mapped to Local Area Network, the test failover machine must have an adapter named Local Area Network. If you have Local Area Network mapped to Network 2, and the test failover machine does not have an adapter named Network 2, test failover will fail.
 - The source, target, and protection job will remain online and uninterrupted during the test.
 - During the test, any scheduled snapshots for the protection job will be deferred until after the test server is online.
 - The test will be performed using the test failover settings configured during job creation.
 - The test failover will take a snapshot of the current data on the target and mirror the data from the snapshot to the test failover machine using the same mirroring options as the protection job.
 - Once the mirror is complete, the test failover machine is rebooted automatically to finalize the test failover process.
 - The test failover machine will maintain its own identity during the test.
 - The test failover machine will maintain its own networking which keeps it isolated from the rest of the network in order to avoid network conflicts and redirecting clients. Applications or functionality that relies on the source networking may not be fully testable with the test machine networking.
 - When using the test failover machine, log in with local credentials to avoid trust relationship issues.

- When you are finished with your test, undo it.
 - When you undo a test failover, the snapshot 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, including your test failover snapshot. (Performing a live failover or failing over to a snapshot will automatically undo any test in progress.)
- **Failover to a snapshot**—Select this option to initiate a full, live failover without using the current data on the target. Instead, select a snapshot and the data on the target will be reverted to that snapshot. This option will not be available if there are no snapshots on the target. To help you understand what snapshots are available, the **Type** indicates the kind of snapshot.
 - **Scheduled**—This snapshot was taken as part of a periodic snapshot.
 - **Deferred**—This snapshot was taken as part of a periodic snapshot, although it did not occur at the specified interval because the job between the source and target was not in a good state.
 - **Manual**—This snapshot was taken manually by a user.
3. Select how you want to handle the data in the target queue.
- **Apply data in target queues before failover or cutover**—All of the data in the target queue will be applied before failover begins. The advantage to this option is that all of the data that the target has received will be applied before failover begins. The disadvantage to this option is depending on the amount of data in queue, the amount of time to apply all of the data could be lengthy.
 - **Discard data in the target queues and failover or cutover immediately**—All of the data in the target queue will be discarded and failover will begin immediately. The advantage to this option is that failover will occur immediately. The disadvantage is that any data in the target queue will be lost.
 - **Revert to last good snapshot if target data state is bad**—If the target data is in a bad state, Carbonite Availability will automatically revert to the last good Carbonite Availability snapshot before failover begins. If the target data is in a good state, Carbonite Availability will not revert the target data. Instead, Carbonite Availability will apply the data in the target queue and then failover. The advantage to this option is that good data on the target is guaranteed to be used. The disadvantage is that if the target data state is bad, you will lose any data between the last good snapshot and the failure.
4. When you are ready to begin failover, click **Failover**.



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

Right before failover occurs, Carbonite Availability 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 Availability 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 failed over because it appears to be a hardware device.

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

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

After a failover 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.

-
5. If you performed a test failover, you can undo it by selecting **Undo Failover or Cutover** in the toolbar. Confirm the undo process when prompted. If configured, the snapshots used for the test failover will be deleted.
-



If you performed a live failover of a full server job but did not enable reverse protection when you configured the job, you will lose the activated target license. To workaround this situation, uninstall Carbonite Availability on the failed over source (currently on the target hardware) using the Carbonite Replication Console or Windows uninstall applet, and the uninstall process will deactivate both the source and target licenses. You can then reinstall Carbonite Availability on the failed over source (currently on the target hardware) and reactivate the source's original license.

Reversing full server jobs

After a full server failover, the source is running on your original target hardware and your target no longer exists. That means the source and target hardware now share the same identity, which is the source identity.



If you did not enable reverse protection, your source is a domain controller, or if you have to rebuild your source, you will have to reverse your protection manually. See *Reversing full server jobs manually* on page 198.

1. Fix the issue that caused your original source server to fail.
2. Connect the original source server to the network.
3. Make sure the production NIC on your original source is online. If the NIC is disabled or unplugged, you will experience problems with the reverse. Make sure you continue to access the servers through the reserved IP addresses, but you can disregard any IP address conflicts for the primary NIC. Since the new source (running on the original target hardware) already has the source's address assigned to it, Windows will automatically assign a different address to the original source.
4. On the **Jobs** page, highlight the job that you want to reverse. If the job is not listed, you may need to add your servers to your console again. Use the reserved IP addresses and local credentials.
5. Highlight the job you want to reverse and click **Reverse** in the toolbar. During the reverse process, you will see various states for the job. During the **Restoring** process, the target identity is being established on the original source hardware. During the **Synchronizing** process, protection is being established from the source (on the original target hardware) to the target (on the original source hardware). The reverse protection is also established in the opposite direction. When the reverse process is complete, the target (on the original source hardware) will reboot. At this point, your source is still running on your original target hardware with the source name, but the original source hardware now has the target identity.
6. To go back to your original hardware, highlight the job and click **Failover**. The source identity will now be applied to the target (on the original source hardware), and the target identity will again be gone. Both servers will have the source identity.
7. To bring back the target identity, highlight the job and click **Reverse**. The same process as above will be repeated, but on the opposite servers. When the reverse is complete, you will be back to your original identities on the original hardware.



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

Right before failover occurs, Carbonite Availability 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 Availability 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 failed over because it appears to be a hardware device.

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

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

After a failover 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.

Reversing full server jobs manually

If you did not enable reverse protection, your source is a domain controller, or if you have to rebuild your source, you have two options after a failover. You can continue running from the failed over server indefinitely. This server is your source (running on the original target hardware) and you can protect it to a new target. Your other option is to go back to the original hardware. Without reverse protection, you have to complete this process manually, which can be difficult.

Preparation of your original source hardware or a new server is key to this manual process. The type of preparation required will depend on the role of the original source server, the applications that were used on the original server, whether the original source was a physical or virtual server, and the failure or event that occurred.

- **Server role**—If your original source was a domain controller, a Cluster Service server, or a Certificate Service server, you will have to reinstall Windows. The utility required to reuse a server cannot be used on these types of servers. Start with *1A. Preparing a new server by reinstalling Windows* on page 199 and then continue with the remaining instructions.
- **Applications**—If your original source was running a name-specific application, like Exchange, you should reinstall Windows. Start with *1A. Preparing a new server by reinstalling Windows* on page 199 and then continue with the remaining instructions.
- **Physical servers**—If your original source was a physical server, your preparation method will depend on if you experienced a catastrophic or non-catastrophic failure.
 - **Catastrophic failure**—If your original hardware is unusable or the failure will require you to reinstall Windows, start with *1A. Preparing a new server by reinstalling Windows* on page 199 and then continue with the remaining instructions.
 - **Non-catastrophic failure**—If the failure did not damage the server or the operating system and you want to reuse the server, start with *1B. Reusing your original source hardware* on page 199 and then continue with the remaining instructions.
- **Virtual servers**—If your original source was a virtual server, your preparation method will depend on if you want to create a new virtual guest or reuse the existing one.
 - **New virtual guest**—If your original guest is unusable, start with *1A. Preparing a new server by reinstalling Windows* on page 199 and then continue with the remaining instructions.



If possible, you can attach any virtual hard disks that survived the failure event to a new virtual guest. Reusing any undamaged disks will decrease the time required to restore data because you can use a difference mirror.

As an alternative to manually creating a new virtual guest, you can let Carbonite Availability automatically provision (create) the new virtual guest for you. If you choose this option, you will need to use the instructions *Creating a full server to ESX job* on page 315 or *Creating a full server to Hyper-V job* on page 261 instead of the instructions in this section.

-
- **Reusing virtual guest**—If the failure did not damage the virtual guest and you want to reuse it, start with *1B. Reusing your original source hardware* on page 199 and then continue with the remaining instructions.

1A. Preparing a new server by reinstalling Windows

1. Install or reinstall Windows on your physical or virtual server using unique, temporary server information. See your Windows documentation for details on installing the operating system.
2. After the operating system installation is complete, install Carbonite Availability using the license key from your original target.
3. After Carbonite Availability is installed and activated, continue with 2. *Mirroring and replicating from the source to the original source hardware or new server and failing over* on page 200.

1B. Reusing your original source hardware

1. Disconnect the original source hardware from the network. For a physical server, you may want to disconnect the network cable. For a virtual server, remove it from the network using your virtual console. You must make sure the original source is completely disconnected before proceeding.
2. After the original source hardware is disconnected from the network, remove the target server identity from Active Directory. You should remove the target's original identity, not the identity of the source which the original target hardware now holds.
3. Keeping the original source hardware disconnected from the network, reboot it and login as the local administrator.
4. Stop all application services on the original source hardware and set them to manual.
5. If you failed over the source IP address, create a new unique IP addresses. See your Windows documentation for details on modifying IP addresses.
6. Modify the original source hardware identity by placing the server into a workgroup. Make sure you reboot when prompted, continuing to keep the server disconnected from the network. See your Windows documentation for details on placing a server into a workgroup.
7. After the reboot, login as the local administrator.
8. Using the Carbonite Replication Console, remove and reinsert the original source server into the server list on the **Servers** page.
9. Double-click on the server in the server list to view the server details page, and then click on the **Edit server properties** link.
10. Under the **Licensing** section, enter the license key from the original target server and click **Add**. If the original source server license key is listed, remove it from the **Current license keys** list. Click **OK** to return to the **Servers** page.
11. Run the Microsoft Sysprep utility to modify SIDs (security identifiers) and the server name. If desired, you can use the original target server name when the utility prompts for a server name. See the Microsoft web site for details on the Sysprep utility.



If the Sysprep utility does not force you to choose a new computer name, you will need to complete the following additional steps.

1. Finish the Sysprep process.
2. Reboot the server and login as the local administrator.
3. Rename the computer manually and reboot when prompted.

The server must be given a new name either via Sysprep or manually after Sysprep has completed before you can proceed.

12. Connect the server to the network and continue with *2. Mirroring and replicating from the source to the original source hardware or new server and failing over* on page 200.

2. Mirroring and replicating from the source to the original source hardware or new server and failing over

1. Using the Carbonite Replication Console, delete the original job from the **Jobs** page.
2. Establish full server protection from your source to the original source hardware or new server. See *Creating a full server job* on page 149. Right-click your source (running on the original target hardware) on the **Servers** page and select **Protect**. Select the same data for protection and specify your original source hardware or new server that you built or modified in step 1A or 1B above on the **Select Target Server** page. Use the options that you used when protecting the source initially, although you can select different settings for snapshots, compression, and so on.
3. Once you have established full server protection, data will be mirrored from the source (on the original target hardware) to the target (on the original source hardware or your new server). Replication will keep the target up-to-date with the changes end-users are continuing to make on the source. Monitor the progress of the job. See *Managing and controlling full server jobs* on page 173.
4. Once the mirror is complete, determine when you want to perform failover. This will require downtime, typically between 15 and 30 minutes depending on LAN or WAN configurations and server processing capabilities.
5. Using the Carbonite Replication Console, perform failover using live data or a snapshot, as desired. See *Failing over full server jobs* on page 192.
6. Monitor the progress. After the target reboots, the target will no longer exist, since it will become the source.

After the reboot, users and other servers can resume normal operations after DNS/IP updates have been propagated to them.

If desired, you can re-establish protection again for this source so that you are prepared for the next emergency.



If you want to reuse the same target hardware, you will have to remove the source identity components from that server. You can use either method *1A. Preparing a new server by reinstalling Windows* on page 199 or *1B. Reusing your original source hardware* on page 199 to prepare a new target.

Chapter 7 SQL protection

Create a SQL job when you have Microsoft SQL Server and want application-level protection.

- *SQL requirements* on page 202—SQL protection includes specific requirements for this type of protection.
- *Creating a SQL job* on page 209—This section includes step-by-step instructions for creating a SQL job.
- *Managing and controlling SQL jobs* on page 232—You can view status information about your SQL jobs and learn how to control these jobs.
- *Failing over SQL jobs* on page 251—Use this section when a failover condition has been met or if you want to failover manually.
- *Restoring then failing back SQL jobs* on page 253—Use this section when you are ready to restore and failback.



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

SQL requirements

Use these requirements for SQL protection.

- **SQL versions**—Carbonite Availability can protect Microsoft SQL Server or Express 2005, 2008, 2008 R2, 2012, 2014, 2016, 2017, or 2019.
- **Operating system**—The following operating systems are supported for SQL jobs.
 - Windows 2019 and Server Core 2019
 - 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 and 2019 support are for the primary operating system features available in Windows 2012. New operating system features specific to Windows 2016 and 2019, such as Nano Server, Windows Containers, and so on, are not supported.

Due to VMware limitations, Windows 2016 and 2019 are not supported on ESX 5.5.

DNS updates are not supported for Server Core servers.

- **SQL and network configuration**—The following requirements and limitations apply to your SQL server and network configuration.
 - All Microsoft best practices should be used for all versions of SQL.
 - You should use the same version, service pack, and architecture (32-bit or 64-bit) of SQL Server on both the source and target servers.
 - SQL 2012 and later are the only supported SQL versions if you are using Windows 2012 in a cluster configuration.
 - SQL 2016 is only supported on Windows 2012 and later.
 - SQL 2017 is only supported on Windows 2012 R2 and later.
 - SQL AlwaysOn Availability Groups (also known as AlwaysOn AG or AAG) is only supported with SQL Server 2016 and SQL 2017 and Windows 2016 in a cluster to standalone environment. No other configurations are currently supported. SQL AlwaysOn Failover Clustering Instances (also known as AlwaysOn FCI or AFCI) is not supported.
 - If you are using SQL Express 2008, you will need to enable and start the SQL Browser Service. By default, this service is set to disabled. Additionally, you will need to enable TCP/IP to accept remote connections. To do this, launch the SQL Server Configuration Manager. Expand **SQL Server Network Configuration**, and under **Protocols** enable **TCP/IP**.
 - If you are using SQL Express 2005, you will need to enable named pipes and TCP/IP to accept remote connections. To do this, launch the SQL Server Configuration Manager. Expand **SQL Server 2005 Network Configuration**, and under **Protocols** enable **Named Pipes** and **TCP/IP**.
 - The SQL program files must be installed in the same location on the source and target.

- The drive letter(s) where SQL stores its data on the source must be the same on the target.
 - The source and target servers must have named instances with the same name installed prior to configuring protection.
 - Single-label DNS domain names (those without a suffix such as .com, .corp, .net) are not supported.
 - In environments where the FIPS security policy is enabled, you must use impersonation, which requires the following.
 - The user running the Carbonite Replication Console must have all appropriate rights to update the domain (that is, only impersonation is supported).
 - You must manually verify DNS rights by running the DFO utility with the /test parameter.
 - If your source and target are in a domain, they should be in the same domain.
 - If your source and target are in a workgroup, make sure the source server's NIC does not register the IP addresses with DNS.
 - If you are using a domain service account that is not in the domain or local Administrators security group, the replicated databases will not mount on the target after failover or on the source after restore because of security restrictions at the file system level. You need to place the SQL 2005 service account in the local Administrators group on the source and target because of the way these versions use local groups for NTFS permissions.
 - You may want to exclude the tempdb database to reduce mirroring and replication traffic. See *Application optimizations* on page 395.
 - Transparent Data Encryption (TDE) is supported for SQL 2008 or later, however the SQL service must be running with the same service account on the source and target.
 - You should use a domain user account or a built-in account (such as NetworkService, LocalService or NT Service) as the Windows Service Account for SQL. See [Setting Up Windows Service Accounts](#) on the Microsoft web site for more information. You should include this account in the local Administrators group on your source and target to ensure that the appropriate permissions for the replicated databases and log files will be available after failover and failback.
 - **File system**—Carbonite Availability supports the NTFS file system. On Windows 2016 and later, 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 21.
 - **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 Availability program files.
-
- **Server name**—Carbonite Availability includes Unicode file system support, but your server name must still be in ASCII format. Additionally, all Carbonite Availability servers must have a unique server name.
-



If you need to rename a server that already has a Carbonite Availability license applied to it, you should deactivate that license before changing the server name. That includes



rebuilding a server or changing the case (capitalization) of the server name (upper or lower case or any combination of case). If you have already rebuilt the server or changed the server name or case, you will have to perform a host-transfer to continue using that license. See the *Carbonite Availability and Carbonite Migrate Installation, Licensing, and Activation* document for complete details.

- **Time**—The clock on your Carbonite Availability servers must be within a few minutes of each other, relative to UTC. Large time skews (more than five minutes) will cause Carbonite Availability 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 console 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 Availability 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 Availability.
- **NAT support**—Carbonite Availability 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 Availability 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 Availability 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 Availability installation program will automatically attempt to configure ports 6320, 6325, and 6326 for Carbonite Availability. 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 367 for instructions on handling firewalls in your environment.

- **Windows Management Instrumentation (WMI)**—Carbonite Availability is dependent on the WMI service. If you do not use this service in your environment, contact technical support.
- **Snapshots**—You can take and failover to Carbonite Availability snapshots using a SQL job.

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

- **Snapshot location**—Snapshots are taken at the volume level and stored on the target. For example, if your job is protecting D:\data and E:\files, the snapshot will contain all of the data on both the D: and E: volumes. If your job is only protecting D:\data (E:\files exists but is not included in the job), the snapshot will only contain the D: volume. Make sure you have enough space on your target for snapshots.
- **Carbonite Availability installation location**—In order to enable Carbonite Availability snapshots, Carbonite Availability must be installed on the system drive. If Carbonite Availability is not installed on the system drive, snapshots will be disabled when enabling protection.
- **Server IP address**—If you have specified an IP address as the source server name, but that IP address is not the server's primary IP address, you will have issues with snapshot functionality. If you need to use snapshots, use the source's primary IP address or its name.
- **Snapshot limitations**—Sometimes taking a snapshot may not be possible. For example, there may not be enough disk space to create and store the snapshot, or maybe the target

is too low on memory. If a snapshot fails, an Event message and a Carbonite Availability log message are both created and logged.

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

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

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

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

Volume Shadow Copy snapshots are associated with the volume they belong to. Since Carbonite Availability 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.

- **Clusters**—If you are using a cluster, make sure your cluster meets the following requirements.
 - **Best practices**—You should carefully review Microsoft documentation and resources for properly configuring your cluster before implementing Carbonite Availability on a cluster. The Microsoft TechNet articles [Failover Clusters](#) and [Installing and Upgrading on Cluster Nodes](#) are two resources you can start with. There are many other resources available on the [Microsoft TechNet web site](#).
 - **Networking**—The following networking requirements apply to your cluster.
 - You must have TCP/IP connections between nodes.
 - Multiple networks are recommended to isolate public and private traffic.
 - The private network should be a unique subnet so that Carbonite Availability will not attempt to use an unreachable private network.
 - Your network can contain direct LAN connections or VLAN technology.
 - The source cluster IP must be accessible from the target.
 - **Domain**—The cluster nodes must be members of the same domain.
 - **DNS**—Forward and reverse lookups must be implemented on the primary DNS server for

the cluster name and individual nodes.

- **Carbonite Availability disk queue**—Ensure that the disk queue is not on a Physical Disk resource.
- **Volumes**—The source and target should have identical drive mappings.
- **Owning nodes**—In a cluster configuration, if you add a possible owning node to the protected network name after a job has started, you must stop and restart the job. If you do not, the records for the new node will not be locked. This could cause problems with DNS records if the source cluster nodes are rebooted or the resources are otherwise cycled on the new owning node.
- **Licensing**—Each node in the cluster must have a valid Carbonite Availability license key.
- **Resource registration**—In some cases, the Carbonite Availability cluster resources may not be registered automatically when Carbonite Availability 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.
- **Supported configurations**—The following table identifies the supported configurations for a SQL job.

Server Configuration	Description	Supported	Not Supported
One to one active/standby	You can protect 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 protect 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 protect many source servers to one target server.		X
One to many	You cannot protect a single source to multiple target servers.		X
Chained	You cannot protect 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 protect a single source to itself.		X
Standalone to standalone	Your servers can be in a standalone to standalone configuration.	X	

Server Configuration	Description	Supported	Not Supported
Standalone to cluster	Your servers cannot be in a standalone to cluster configuration.		X
Cluster to standalone	Your servers can be in a cluster to standalone configuration.	X	
Cluster to cluster	Your servers can be in a cluster to cluster configuration.	X	

If you are using a Cluster Shared Volume (CSV), you cannot protect any data, including a virtual machine, residing on the CSV. If you want to protect a CSV virtual machine, you must run Carbonite Availability 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 SQL job

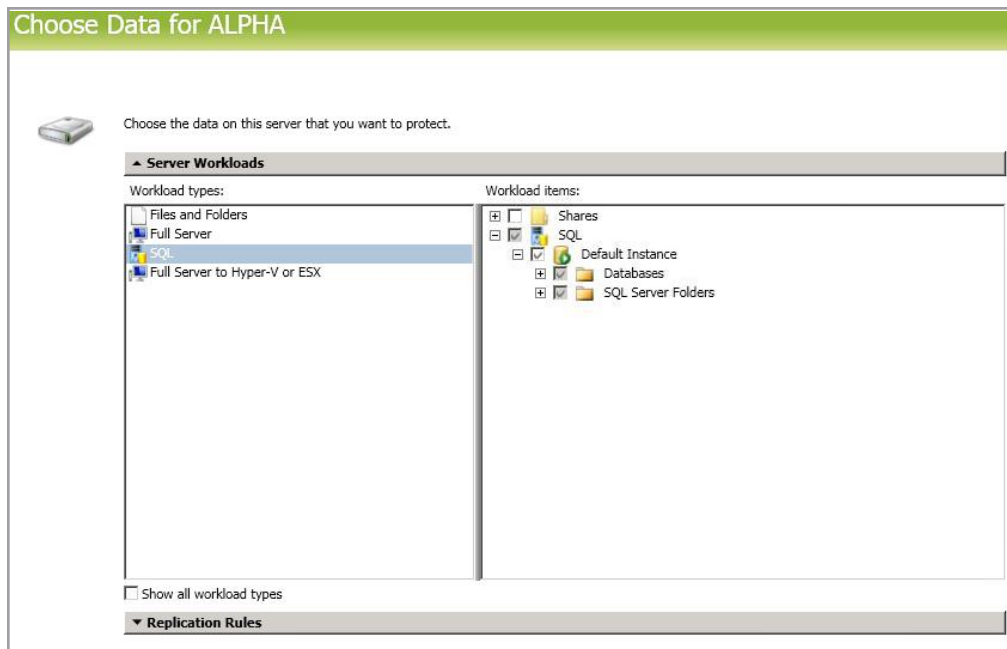
Use these instructions to create a SQL job.

1. Make sure you are logged in to the Carbonite Replication Console as a user with the SQL sysadmin role before you begin the SQL job creation process.
2. From the **Servers** page, right-click the server you want to protect and select **Protect**. You can also highlight a server, click **Create a New Job** in the toolbar, then select **Protect**.
3. Choose the type of workload that you want to protect. Under **Server Workloads**, in the **Workload types** pane, select **SQL Server**. In the **Workload items** pane, Carbonite Availability will automatically select the entire SQL program and data files. If desired, you can exclude user databases from protection, but the system databases (except for tempdb) are required and cannot be excluded.



New shares within the original replication rules created after the job is created will be included on the target. However, new shares created outside of the original replication rules after the job is created will not be included on the target. For example, if C:\a is originally protected and a new share for C:\a\b is created, that share will be included on the target. However if C:\b is created, you must modify the replication rules to make sure the new share is included on the target.

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.



4. If you want to select other files and folders to include in your protection, click the **Replication Rules** heading and expand the volumes under **Folders**.

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

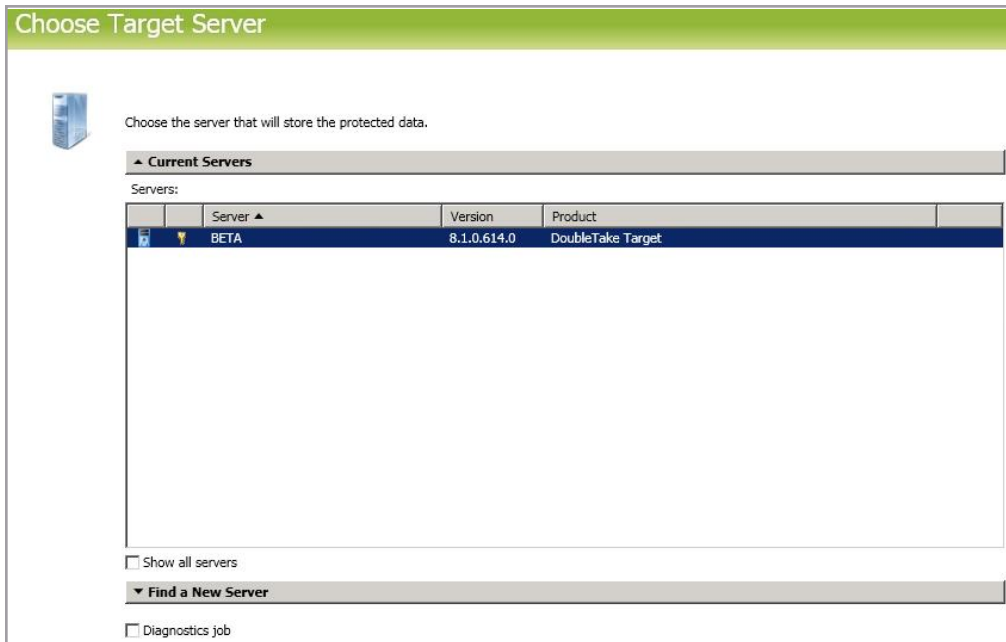
You can also enter wildcard rules, however you should do so carefully. Rules are applied to files that are closest in the directory tree to them. If you have rules that include multiple folders, an exclusion rule with a wild card will need to be added for each folder that it needs applied to. For example,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

If you need to remove a rule, highlight it in the list at the bottom and click **Remove Rule**. Be careful when removing rules. Carbonite Availability 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.

5. Click **Next** to continue.
6. Choose your target server. This is the server that will store the replica SQL Server from the source.



- **Current Servers**—This list contains the servers currently available in your console session. Servers that are not licensed for the workflow you have selected and those not applicable to the workload type you have selected will be filtered out of the list. Select your target server from the list. 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. Additionally, the user must meet the following credentials requirements.

- The account must be assigned the sysadmin role on the SQL server in order to query and administer SQL, and the SQL service startup account should be a domain account.
 - The account must be a member of the Domain Admins group. If your security policies do not allow use of this group, see *DNS* on page 375 and use the instructions under the Carbonite Availability DFO utility to use a non-Domain Admins account.
-

7. 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.

8. You have many options available for your SQL job. Configure those options that are applicable to your environment.

Go to each page identified below to see the options available for that section of the **Set Options** page. After you have configured your options, continue with the next step on page 231.

- *General* on page 212
- *Failover Monitor* on page 213
- *Failover Options* on page 216
- *Test Failover Scripts* on page 218
- *Failover Identity* on page 219
- *Mirror, Verify & Orphaned Files* on page 221
- *Network Route* on page 224
- *Target Paths* on page 225
- *Snapshots* on page 226

- *Compression* on page 227
- *Bandwidth* on page 228
- *Scripts* on page 230

General



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

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

Failover Monitor

Failover Monitor

Total time to failure: 00:05:00

Consecutive failures: 5

Monitor on this interval: 00:00:05

Network monitoring

Monitor these addresses:

	Source IP Address
<input checked="" type="checkbox"/>	112.42.74.29
<input checked="" type="checkbox"/>	10.10.10.29

Monitoring method: Network service

Failover trigger: All monitored IP addresses fail

Service monitoring

Services to monitor:

SQL Server (MSSQLSERVER)

Add... Remove

Attempt to restart this service after each failure

Custom script monitoring

Script file: Browse...

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

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

- **Consecutive failures**—Specify how many attempts the target will make to contact the source before the source is considered failed. For example, if you have this option set to 20, and your source fails to respond to the target 20 times in a row, this will be considered a failure.
- **Monitor on this interval**—Specify, in hours:minutes:seconds, how long to wait between

attempts to contact the source to confirm it is online. This means that after a response (success or failure) is received from the source, Carbonite Availability will wait the specified interval time before contacting the source again. If you set the interval to 00:00:00, then a new check will be initiated immediately after the response is received.

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

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

- **Network monitoring**—With this option, the target will monitor the source using a network ping.
 - **Monitor these addresses**—Select each **Source IP Address** that you want the target to monitor. If you are protecting a cluster, you are limited to the IP addresses in the cluster group that you are protecting. If you are using a NAT environment, do not select a private IP address on the source because the target cannot reach the source's private address, thus causing an immediate failure. Also for NAT environments, you will see an additional field for the **Replication Service port**. This gives you the ability to specify the port number to be used with the address, allowing the target to monitor the source through a router.
 - **Monitoring method**—At this time, **Network service** is the only option. Failover monitoring will test source availability by using an ICMP ping to confirm the route is active.
 - **Failover trigger**—If you are monitoring multiple IP addresses, specify when you want a failover condition to be triggered.
 - **One monitored IP address fails**—A failover condition will be triggered when any one of the monitored IP addresses fails. If each IP address is on a different subnet, you may want to trigger failover after one fails.
 - **All monitored IP addresses fail**—A failover condition will be triggered when all monitored IP addresses fail. If there are multiple, redundant paths to a server, losing one probably means an isolated network problem and you should wait for all IP addresses to fail.
- **Service monitoring**—This option is only available in standalone environments. It is not available for clustered source servers. With this option, the target will monitor specific services on the source by confirming that they are running. Multiple services in the list will be checked in parallel. A failover condition is met when one of the monitored services fails the check. Click **Add** and select the service that you want to monitor. Repeat this step for additional services that you want to monitor. If you want to remove a service from the **Services to monitor** list, highlight it and click **Remove**.
 - **Attempt to restart this service after each failure**—When this option is enabled, if a service fails the monitor check, Carbonite Availability will attempt to restart it. During this restart period, Carbonite Availability will not check the other services, to avoid any false failures while the one service is attempting to be restarted. If the service cannot be restarted, Carbonite Availability will consider this a failure.

- **Custom script monitoring**—This option is only available in standalone environments. It is not available for clustered source servers. With this option, you can use your own custom script to monitor the source for a failure. You can use any standard script such as PowerShell, VB, batch files, and so on. Your script must return a code of 0 upon success, unless you are using a PowerShell script, in which case you should be using an exit value instead of a return value. Any other code will indicate a failure condition has been met. Your script will not be interactive, so ensure that it does not require any user interaction to execute successfully.
-



The network monitoring test is performed independently, while the service and custom monitoring tests are performed in parallel. Therefore, if you are using network monitoring with service and/or custom monitoring, and the network monitor fails, the service and custom monitoring will be skipped. If the network monitor is successful, service or custom monitoring will be performed. Any failure between these two tests will be considered a failure, therefore, both tests must complete and pass for the test to be considered successful.

Failover Options

Failover Options

Wait for user to initiate failover
 Failover shares

Active Directory
 Failover host name
 Failback host name
Active Directory Credentials...

Target scripts

Pre-failover script: [] Arguments: []
 Delay failover until script completes

Post-failover script: [] Arguments: []

Pre-failback script: [] Arguments: []
 Delay failback until script completes

Post-failback script: [] Arguments: []

Source scripts

Post-failback script: [] Arguments: []

- **Wait for user to initiate failover**—The failover process can wait for you to initiate it, allowing you to control when failover occurs. When a failure occurs, the job will wait in **Failover Condition Met** for you to manually initiate the failover process. Disable this option if you want failover to occur immediately when a failure occurs.
- **Failover shares**—Select this option to failover shares to the target. Only the shares that you selected on the **Choose Data** page will be protected and failed over.



Share failover only occurs for standard Windows file system shares. Other shares must be configured for failover through the failover scripts or created manually on the target. See *Macintosh shares* on page 385 or *NFS Shares* on page 386 for more information.

If your target is a standalone server, Windows share information is automatically updated on the target every 30 seconds. Since shares are not failed over if your target is a cluster, you will need to manually update shares on a cluster target. You will need to repeat the initial cluster configuration steps, although you can skip giving the target cluster account full control and creating the target Client Access Point, because those steps were already completed during the initial cluster configuration.

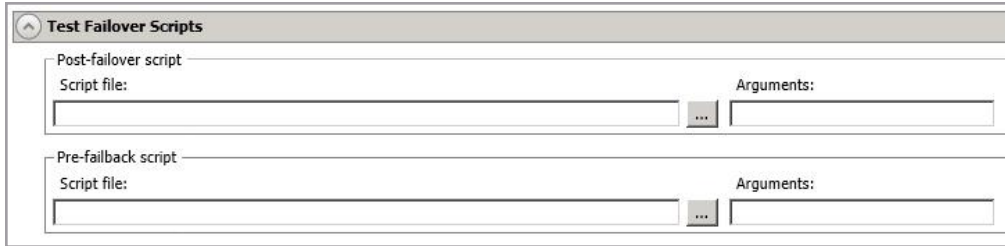
- **Failover host name**—If desired, you can failover the source server's host name. This will automatically remove the host SPN (Service Principal Name) from Active Directory on the source and add it to Active Directory on the target. If you are using two different accounts

for the SQL service login on the source and target, you should failover the SPNs. If you are using the same domain account for the SQL service login on the source and target, you do not need to failover the SPNs.

- **Failback host name**—This option returns the host SPN on the source and target back to their original settings on failback. If you are using Active Directory, enable this option or you may experience problems with failback.
- **Active Directory Credentials**—If you are failing over and/or failing back the host name, you need to specify a user that has update privileges within Active Directory. Click **Active Directory Credentials** and identify a user and the associated password that has privileges to create and delete SPNs. The username must be in the format `fully_qualified_domain\user`, and the account password cannot be blank.
- **Scripts**—You can customize failover and failback by running scripts on the source and target. Scripts may contain any valid Windows command, executable, or batch file. The scripts are processed using the same account running the Double-Take Management service, unless you have identified a specific account through the server's properties. See *Script credentials* on page 67. Examples of functions specified in scripts include stopping services on the target before failover because they may not be necessary while the target is standing in for the source, stopping services on the target that need to be restarted with the source's machine name and/or IP address, starting services or loading applications that are in an idle, standby mode waiting for failover to occur, notifying the administrator before and after failover or failback occurs, stopping services on the target after failback because they are no longer needed, stopping services on the target that need to be restarted with the target machine's original name and/or IP address, and so on. There are four types of failover and failback scripts that run on the target and one failback script that runs on the source. There are also test scripts you can specify that will run when you are testing the failover process.
 - **Pre-failover script**—This script runs on the target at the beginning of the failover process. Specify the full path and name of the script file.
 - **Post-failover script**—This script runs on the target at the end of the failover process. Specify the full path and name of the script file.
 - **Pre-failback script**—This script runs on the target at the beginning of the failback process. Specify the full path and name of the script file.
 - **Post-failback script**—This script runs on the target or source at the end of the failback process. Specify the full path and name of the script file.
 - **Arguments**—Specify a comma-separated list of valid arguments required to execute the script.
 - **Delay until script completes**—Enable this option if you want to delay the failover or failback process until the associated script has completed. If you select this option, make sure your script handles errors, otherwise the failover or failback process may never complete if the process is waiting on a script that cannot complete.

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

Test Failover Scripts



This section is not applicable to clustered environments.

When you failover, you will have three choices. You can failover to live data, failover to data from a snapshot, or perform a test failover. Any scripts you specify in this section are only used during a test failover.

- **Post-Failover Script**—This script runs on the target at the end of the failover process. Specify the full path and name of the script file.
- **Pre-Failback Script** —This script runs on the target at the beginning of the undo failover process. Specify the full path and name of the script file.
- **Arguments**—Specify a comma-separated list of valid arguments required to execute the script.

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)

Failover server name (NetBIOS)

Update DNS server

DNS Options

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

These DNS servers will be updated during failover:

112.42.48.9

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

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

Update TTL (seconds):
300



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

- **Retain target network configuration**—The target will retain all of its original IP addresses.
 - **Failover server name**—Select this option if you want to failover the NetBIOS name. This option will not be available for clustered targets.
 - **Update DNS server**—Specify if you want Carbonite Availability to update your DNS server on failover. If DNS updates are made, the DNS records will be locked during failover. Be sure and review the job requirements for updating DNS.



Because of changes to when resources in Windows 2012 clusters can come online related to DNS record locking, DNS updates are required for Windows 2012 cluster configurations.

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.

Expand the **DNS Options** section to configure how the updates will be made. The DNS information will be discovered and displayed. If your servers are in a

workgroup, you must provide the DNS credentials before the DNS information can be discovered and displayed.

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



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

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

Mirror, Verify & Orphaned Files

Mirror, 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 Availability 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 the mirror time 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 Availability 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 the mirror time 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 Availability 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 speed. 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 Availability 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 speed.

- **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 Availability 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 Availability 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 Availability 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.

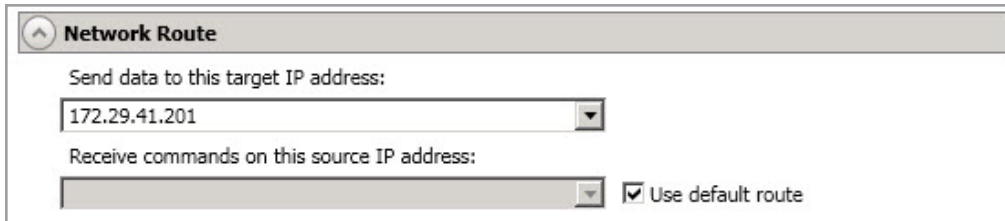
If delete orphaned files is enabled, carefully review any replication rules that use wildcard definitions. If you have specified wildcards to be excluded from protection, files matching those wildcards will also be excluded from orphaned file processing and will not be deleted from the target. However, 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 63 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 Availability 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 Availability log, and during difference mirrors, verifications, and restorations, all orphaned file processing messages are logged to the Carbonite Availability log. The orphaned file log is located in the **Logging folder** specified for the source. See *Log file properties* on page 68 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 this target IP address:

172.29.41.201

Receive commands on this source IP address:

Use default route



If your target is a cluster, you will see a table presented where you can select the target route, rather than the single target route field as shown above. The field functions the same as described below.

- **Send data to this target IP address**—By default, Carbonite Availability will select an IP address on the target 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 Availability traffic. For example, you can separate regular network traffic and Carbonite Availability 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.



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

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

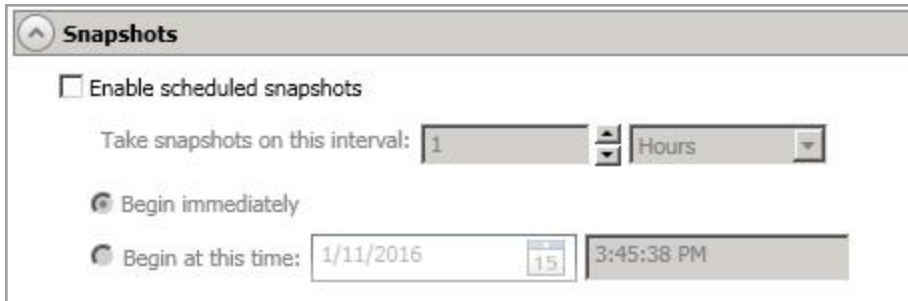
- **Receive commands on this source IP address**—By default, Carbonite Availability will select an IP address on the source to receive commands and requests for status from the target. This is communication from the Double-Take Management Service. If desired, specify an alternate route on the source that the commands and requests will be transmitted to. This allows you to select a different route for Carbonite Availability management communication. You can also manually enter a public IP address (which is the public IP address of the source server's router) if you are using a NAT environment.
- **Use default route**—Select this option to disable the drop-down list that allows you to select the route from the target server. When this option is enabled, the default route will automatically be used.

Target Paths



Block target paths upon connection allows you to block writing to the replica source data located on the target. This keeps the data from being changed outside of Carbonite Availability processing. Any target paths that are blocked will be unblocked automatically during the failover process so that users can modify data after failover. During restoration, the paths are automatically blocked again. If you failover and failback without performing a restoration, the target paths will remain unblocked.

Snapshots



Snapshots

Enable scheduled snapshots

Take snapshots on this interval: 1 Hours

Begin immediately

Begin at this time: 1/11/2016 3:45:38 PM

A snapshot is an image of the source replica data on the target taken at a single point in time. You can view the snapshots in VSS and recover any files or folders desired. You can also failover to a snapshot.

Turn on **Enable scheduled snapshots** if you want Carbonite Availability to take snapshots automatically at set intervals.


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



See *Managing snapshots* on page 76 for details on taking manual snapshots and deleting snapshots.

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

Compression



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

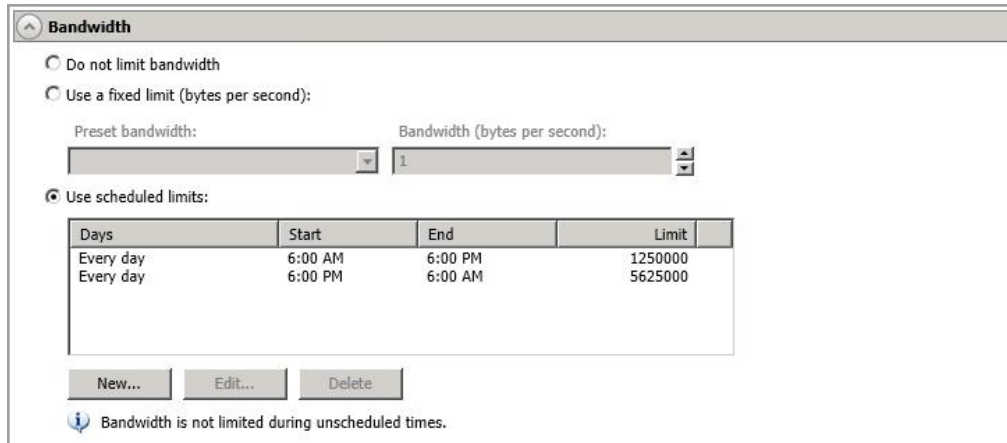
Keep in mind that the process of compressing data impacts processor usage on the source. If you notice an impact on performance while compression is enabled in your environment, either adjust to a lower level of compression, or leave compression disabled. Use the following guidelines to determine whether 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 Availability 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 Availability data transmissions. When a bandwidth limit is specified, Carbonite Availability never exceeds that allotted amount. The bandwidth not in use by Carbonite Availability is available for all other network traffic.



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

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

- **Do not limit bandwidth**—Carbonite Availability will transmit data using 100% bandwidth availability.
- **Use a fixed limit**—Carbonite Availability 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 Availability 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 for configuring scripts:

- 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.

If you are using Windows 2008 R2 with Service Pack 1, you can customize mirroring by running scripts on the target at pre-defined points in the mirroring process. Scripts are not functional for Windows 2012 and later.

Scripts may contain any valid Windows command, executable, or batch file. The scripts are processed using the same account running the Double-Take service, unless you have identified a specific account through the server's properties. See *Script credentials* on page 67. 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).

9. Click **Next** to continue.
10. Carbonite Availability 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 Availability correct the problem for you. For those errors that Carbonite Availability 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 Availability log files.

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



Jobs in a NAT environment may take longer to start.

Once a job is created, do not change the name of underlying hardware components used in the job. For example, volume names, network adapter names, or 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 SQL 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 33), then only the jobs associated with the server or target servers in that server group will be displayed in the right pane.



In a cluster environment, when a SQL job is stopped, the DTJobStatus resource (because of resource dependencies) will prevent the SQL application from coming online on the source while the target is unavailable. If you need to override these dependencies in order to bring the SQL application online, you can enable the **Override Dependencies** option on the **Parameters** tab of the DTJobStatus resource in Failover Cluster Manager.

-
- *Overview job information displayed in the top right pane* on page 232
 - *Detailed job information displayed in the bottom right pane* on page 235
 - *Job controls* on page 237

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 Availability 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 a name or IP address of a standalone server, a cluster, or a node. Cluster jobs will be associated with the cluster name and standalone jobs will be associated with a standalone server or a cluster node.

Target Server

The name of the target. This could be a name or IP address of a standalone server, a cluster, or a node. Cluster jobs will be associated with the cluster name and standalone jobs will be associated with a standalone server or a cluster node.

Job Type

Each job type has a unique job type name. This job is a SQL 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.
- **Restoring**—Data is being restored from the target to the source.
- **Unknown**—The console cannot determine the status.

Replication Status

- **Replicating**—Data is being replicated to the target.
- **Ready**—There is no data to replicate.
- **Pending**—Replication is pending.
- **Stopped**—Replication has been stopped.
- **Out of Memory**—Replication memory has been exhausted.

- **Failed**—The Double-Take service is not receiving replication operations from the Carbonite Availability 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.
- **Restore Required**—The data on the source and target do not match because of a failover condition. Restore the data from the target back to the source. If you want to discard the changes on the target, you can remirror to resynchronize the source and target.
- **Snapshot Reverted**—The data on the source and target do not match because a snapshot has been applied on the target. Restore the data from the target back to the source. If you want to discard the changes on the target, you can remirror to resynchronize the source and target.
- **Busy**—The source is low on memory causing a delay in getting the state of the data on the target.
- **Not Loaded**—Carbonite Availability 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 54.

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 54.

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. This is the current time where the console is running.

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 **Edit Job 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.

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 Availability queue was not exhausted during the time the job was paused. If the Carbonite Availability queue was exhausted during the time the job was paused, the job will restart mirroring and replication.

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

Pause

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

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

Stop

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

Take Snapshot

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

Manage Snapshots

Allows you to manage your snapshots by taking and deleting snapshots for the selected job. See *Managing snapshots* on page 76 for more information.

Failover or Cutover

Starts the failover process. See *Failing over SQL jobs* on page 251 for the process and details of failing over a SQL job.

Failback

Starts the failback process. See *Restoring then failing back SQL jobs* on page 253 for the process and details of failing back a SQL job.

You may receive an error when trying to failback to the source cluster if the application's virtual IP address is offline. Verify the source cluster's virtual name and IP address resources are online, as well as the application's virtual IP address resource, and retry failback.

Restore

Starts the restoration process. See *Restoring then failing back SQL jobs* on page 253 for the process and details of restoring a SQL job.

Reverse

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

Undo Failover or Cutover

Cancels a test failover by undoing it. Undo failover does not apply to clustered jobs. See *Failing over SQL jobs* on page 251 for details on undoing a test failover.

View Job Log

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

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 Availability 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 Availability 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 Availability 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 the mirror time 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 Availability 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 the mirror time 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 Availability 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 speed. 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 Availability 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 speed.
- **Calculate size of protected data before mirroring**—Specify if you want Carbonite Availability 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 Availability will transmit data using 100% bandwidth availability.
 - **Fixed bandwidth limit**—Carbonite Availability 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 Availability 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 Availability processing, not the entire server.

While the target is paused, the Carbonite Availability 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 Availability 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.

- **Refresh Status**—Refreshes the job status immediately.

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 33.

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 SQL 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 SQL 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.
- **Restore Required**—The data on the source and target do not match because of a failover condition. Restore the data from the target back to the source. If you want to discard the changes on the target, you can remirror to resynchronize the source and target.
- **Snapshot Reverted**—The data on the source and target do not match because a snapshot has been applied on the target. Restore the data from the target back to the source. If you want to discard the changes on the target, you can remirror to resynchronize the source and target.
- **Busy**—The source is low on memory causing a delay in getting the state of the data on the target.
- **Not Loaded**—Carbonite Availability 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 Availability 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.
- **Restoring**—Data is being restored from the target to the source.
- **Unknown**—The console cannot determine the status.

Mirror percent complete

The percentage of the mirror that has been completed

Mirror remaining

The total number of mirror bytes that are remaining to be sent from the source to the target. This value may be zero if you have enabled **Mirror only changed files when source reboots** on the *Server setup properties* on page 54.

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 54.

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 Availability 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 SQL job

Over time, you may want to confirm that any changes in your network or environment have not impacted your Carbonite Availability 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 Availability 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 Availability correct the problem for you. For those errors that Carbonite Availability 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 SQL job

Use these instructions to edit a SQL 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 Availability monitor security access.)
2. You will see the same options available for your SQL 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 SQL job* on page 209 for details on each job option.



Changing some options may require Carbonite Availability 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 your job is a clustered SQL job, you will not be able to edit the workload or replication rules. However, if you have a standalone SQL job and 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 Availability 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 Availability orphan file detection cannot remove the data for you. Therefore, you have to remove it manually.

4. Click **Next** to continue.
5. Carbonite Availability 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 Availability correct the problem for you. For those errors that Carbonite Availability 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 Availability log files.

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

Viewing a SQL job log

You can view a job log file through the 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.



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

Time	Description
6/22/2017 2:42:52 PM	Hardware IDs as follows: Source = 'bb1e75f-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_6d2bf...
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.

Failing over SQL jobs

When a failover condition has been met, failover will be triggered automatically if you disabled the wait for user option during your failover configuration. If the wait for user before failover option is enabled, you will be notified in the console when a failover condition has been met. At that time, you will need to trigger it manually from the console when you are ready.



If you have paused your target, failover will not start if configured for automatic failover, and it cannot be initiated if configured for manual intervention. You must resume the target before failover will automatically start or before you can manually start it.

1. On the **Jobs** page, highlight the job that you want to failover and click **Failover or Cutover** in the toolbar.
2. Select the type of failover to perform.
 - **Failover to live data**—Select this option to initiate a full, live failover using the current data on the target. The application services will be stopped on the source (if it is online), and they will be started on the target. DNS records will be updated, if applicable. Application requests destined for the source server or its IP addresses are routed to the target.
 - **Perform test failover**—Select this option to perform a test failover using the current data on the target. This option will allow you to confirm SQL on the target is viable for failover. This process will pause the target (the entire target is not paused, just Carbonite Availability processing), take a snapshot of the target (so that you can revert back to the pre-test state after the test is complete), and then start the application services on the target. Success and failure messages will be available in the job log. (Note the application services are still running on the source during the test.) Once the application services are running on the target, you can perform any testing desired on the target server.



The following caveats apply to performing a test failover for SQL.

- Make sure you have enough space on your target to contain the snapshot.
 - During the test, any schedule snapshots for the protection job will be deferred until after the test.
 - If any of your SQL data is stored on the system volume, you will not be able to perform a test failover.
 - Any scripts you specified in the **Test Failover Scripts** section when you created your job will be executed during the test failover.
-
- **Failover to a snapshot**—Select this option to initiate a full, live failover without using the current data on the target. Instead, select a snapshot and the data on the target will be reverted to that snapshot. This option will not be available if there are no snapshots on the target. To help you understand what snapshots are available, the **Type** indicates the kind of snapshot.
 - **Scheduled**—This snapshot was taken as part of a periodic snapshot.
 - **Deferred**—This snapshot was taken as part of a periodic snapshot, although it did

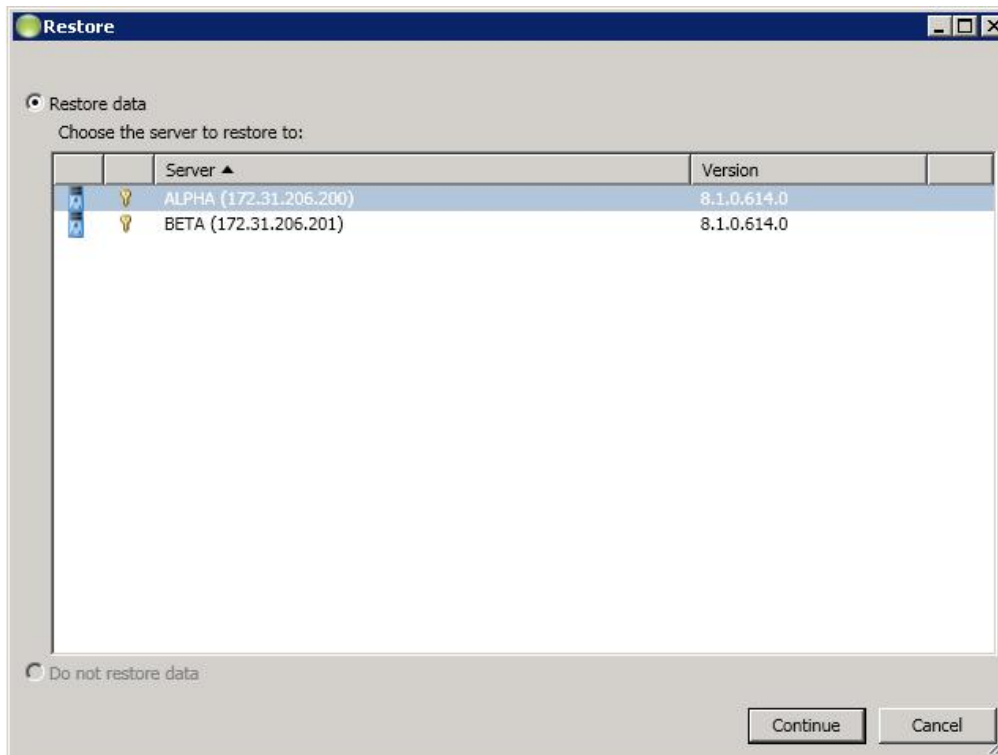
not occur at the specified interval because the job between the source and target was not in a good state.

- **Manual**—This snapshot was taken manually by a user.
3. Select how you want to handle the data in the target queue.
 - **Apply data in target queues before failover or cutover**—All of the data in the target queue will be applied before failover begins. The advantage to this option is that all of the data that the target has received will be applied before failover begins. The disadvantage to this option is depending on the amount of data in queue, the amount of time to apply all of the data could be lengthy.
 - **Discard data in the target queues and failover or cutover immediately**—All of the data in the target queue will be discarded and failover will begin immediately. The advantage to this option is that failover will occur immediately. The disadvantage is that any data in the target queue will be lost.
 - **Revert to last good snapshot if target data state is bad**—If the target data is in a bad state, Carbonite Availability will automatically revert to the last good Carbonite Availability snapshot before failover begins. If the target data is in a good state, Carbonite Availability will not revert the target data. Instead, Carbonite Availability will apply the data in the target queue and then failover. The advantage to this option is that good data on the target is guaranteed to be used. The disadvantage is that if the target data state is bad, you will lose any data between the last good snapshot and the failure.
 4. The last option will vary depending on if your source is clustered.
 - **Leave source services running**—If your source is standalone, select this option if you want the SQL services to remain online instead of being shutdown. This is for live or snapshot failover. It does not apply to test failover.
 - **Leave source cluster resources online**—If your source is clustered, select this option if you want the cluster resources to remain online instead of being taken offline. This is for live or snapshot failover. It does not apply to test failover.
 5. When you are ready to begin failover, click **Failover**.
 6. If you performed a test failover, you can undo it by selecting **Undo Failover or Cutover** in the toolbar. Confirm the undo process when prompted. The application services will be stopped on the target, the pre-test snapshot that was taken will be reverted, and the target will be resumed.

Restoring then failing back SQL jobs

Restoring before failing back allows your users to continue accessing their data on the failed over target, which is standing in for the source, while you perform the restoration process.

1. Resolve the problem(s) on the source that caused it to fail.
2. Bring the source onto the network.
3. On the **Jobs** page, highlight the job and click **Restore**.



4. Confirm **Restore data** is selected, then highlight your source server in the server list.



If you do not want to restore your data, you can select **Do not restore data**. Keep in mind, that any data changes that took place on the target after failover will be lost.

5. Click **Continue** to start the restoration.
6. During the restoration process, the **Activity** will indicate **Restoring** and **Mirror Status** will indicate **In Progress**. When the restoration is complete, the **Mirror Status** will change to **Idle**, and the **Activity** will be **Restored**. At this time, schedule a time for failback. User downtime will begin once failback is started, so select a time that will have minimal disruption on your users.
7. On the **Jobs** page, highlight the job and click **Failback**.
8. In the dialog box, highlight the job that you want to failback and click **Failback**.
9. Check that your source is fully functional and that the source data is in a good state, and then, if desired, you can enable protection again by clicking **Start**.

Chapter 8 Full server to Hyper-V protection

Create a full server to Hyper-V job when you want to protect an entire physical server or virtual machine to a Hyper-V target. There is no reverse protection for this job. You will have to use another full server job type to get back to your original hardware after failover.

- *Full server to Hyper-V requirements* on page 255—Full server to Hyper-V protection includes specific requirements for this type of protection.
- *Creating a full server to Hyper-V job* on page 261—This section includes step-by-step instructions for creating a full server to Hyper-V job.
- *Managing and controlling full server to Hyper-V jobs* on page 283—You can view status information about your full server to Hyper-V jobs and learn how to control these jobs.
- *Failing over full server to Hyper-V jobs* on page 302—Use this section when a failover condition has been met or if you want to failover manually.
- *Reversing protection after failover for full server to Hyper-V jobs* on page 306—Use this section if you need to get your data back to the original hardware.

Full server to Hyper-V requirements

Use these requirements for full server to Hyper-V protection.

- **Source server**—The following operating systems are supported for on the source for full server to Hyper-V jobs.
 - Windows 2019 and Server Core 2019
 - 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 and 2019 support are for the primary operating system features available in Windows 2012. New operating system features specific to Windows 2016 and 2019, such as Nano Server, Windows Containers, and so on, are not supported.

Due to VMware limitations, Windows 2016 and 2019 are not supported on ESX 5.5.

DNS updates are not supported for Server Core servers.

If your source is a Hyper-V server, you will be able to protect it, however the Hyper-V role and features will not be available after failover.

-
- **Target host**—The target host can be any of the following Windows operating systems with the Hyper-V role enabled. The target host can only be used in the combinations described after the operating system list.
 - Windows 2019 and Server Core 2019
 - 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 2008 (non-R2) is not supported for any job type.
 - Hyper-V Server 2008 R2, 2012, 2012 R2, 2016, and 2019



Windows 2008, Hyper-V Server 2008, and Server Core 2008 are not supported.

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

- **2019, 2016, or 2012 R2**—If your host is Windows 2019, 2016, or 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, your source can be Windows 2012 R2 or earlier..
-



If your source is Windows 2016 and you are using ReFS volumes, your target must be Windows 2016. You cannot use a Windows 2019 target because the formatting of ReFS is different between Windows 2016 and Windows 2019.

- **File system**—Carbonite Availability supports the NTFS file system. On Windows 2016 and later, 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 21.
-



Because ReFS is formatted differently on Windows 2016 and Windows 2019, you cannot use a Windows 2016 source with ReFS to a Windows 2019 target.

- **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 Availability program files.
 - **Server name**—Carbonite Availability includes Unicode file system support, but your server name must still be in ASCII format. Additionally, all Carbonite Availability servers must have a unique server name.
-



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

- **Time**—The clock on your Carbonite Availability servers must be within a few minutes of each other, relative to UTC. Large time skews (more than five minutes) will cause Carbonite Availability 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.
 - WAN failover is not supported with IPv6 addresses.
 - If you are using IPv6 on your servers, your console 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 Availability 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 Availability.
- **Network adapters**—Your source can have no more than 12 NICs enabled (eight synthetic and four legacy).
- **NAT support**—Carbonite Availability 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 Availability 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 Availability 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 Availability installation program will automatically attempt to configure ports 6320, 6325, and 6326 for Carbonite Availability. 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 367 for instructions on handling firewalls in your environment.

- **Windows Management Instrumentation (WMI)**—Carbonite Availability is dependent on the WMI service. If you do not use this service in your environment, contact technical support.
- **Snapshots**—You can take and failover to Carbonite Availability snapshots using a full server to Hyper-V job.

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

- **Snapshot location**—Snapshots are taken and stored inside the replica virtual disk, so be sure that you configure your replica virtual machine disks large enough to maintain snapshots.
- **Carbonite Availability installation location**—In order to enable Carbonite Availability snapshots, Carbonite Availability must be installed on the system drive. If Carbonite Availability is not installed on the system drive, snapshots will be disabled when enabling protection.
- **Server IP address**—If you have specified an IP address as the source server name, but that IP address is not the server's primary IP address, you will have issues with snapshot functionality. If you need to use snapshots, use the source's primary IP address or its name.
- **Snapshot limitations**—Sometimes taking a snapshot may not be possible. For example, there may not be enough disk space to create and store the snapshot, or maybe the target is too low on memory. If a snapshot fails, an Event message and a Carbonite Availability log message are both created and logged.

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

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

You must also keep in mind that if you are using extended functionality provided by Volume Shadow Copy, you need to be aware of the impacts that functionality may have on Carbonite Availability. For example, if you change the location where the shadow copies are stored and an error occurs, it may appear to be a Carbonite Availability error when it is

in fact a Volume Shadow Copy error. Be sure and review any events created by the VolSnap driver and check your Volume Shadow Copy documentation for details.

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

Volume Shadow Copy snapshots are associated with the volume they belong to. Since Carbonite Availability 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.

- **Supported configurations**—The following table identifies the supported configurations for a full server to Hyper-V job.

Server to Host Configuration	Description	Supported	Not Supported
One to one active/standby	You can protect a single source to a single target host.	X	
One to one active/active	This configuration (where both the source and target use the same job type to actively replicate to each other) is not supported and not applicable because the target is a hypervisor host.		X
Many to one	You can protect 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.	X	
One to many	You can protect a single source to multiple target hosts. The source is the only server actively replicating data. This will create redundant copies of your source.	X	
Chained	This configuration (where the source replicates to the target and then the target uses the same job type to replicate the source to a final target) is not supported and not applicable because the middle target is a hypervisor host.		X
Single server	You cannot protect a single source to itself.		X

Server to Host Configuration	Description	Supported	Not Supported
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 Availability disk queue— Ensure that the disk queue is on a locally attached, non-cluster volume. • Resource registration—In some cases, the Carbonite Availability cluster resources may not be registered automatically when Carbonite Availability 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	A source cluster or cluster node is not a default supported configuration, but it may be possible with assistance from Professional Services. If you want to use a full server to Hyper-V job in a source cluster environment, 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	Your servers cannot be in a cluster to cluster configuration.		X

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

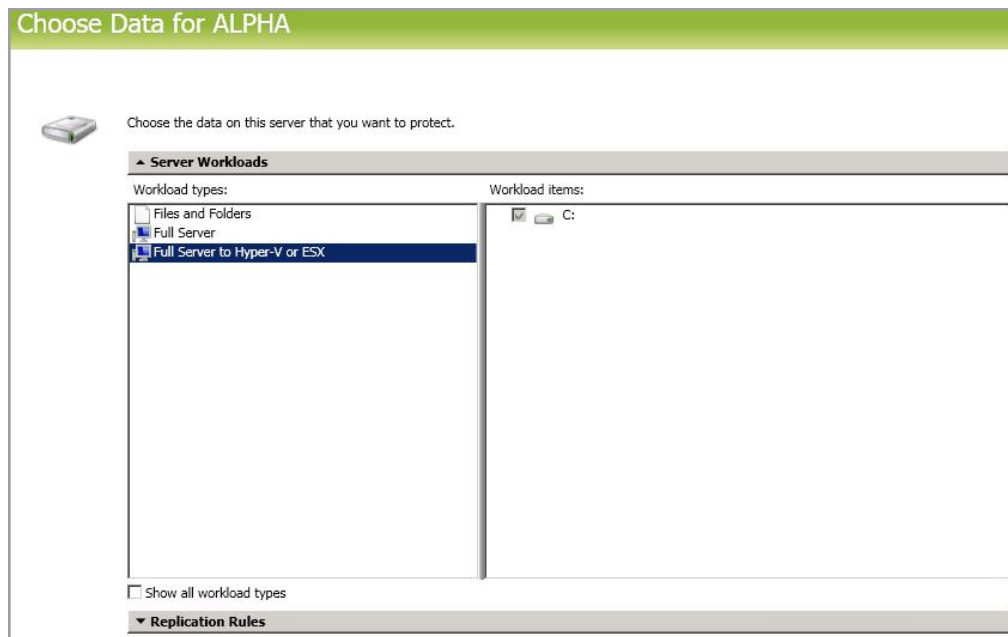
Data can be written to a target CSV.

Creating a full server to Hyper-V job

Use these instructions to create a full server to Hyper-V job.

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

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



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

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

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


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

If you need to remove a rule, highlight it in the list at the bottom and click **Remove Rule**. Be careful when removing rules. Carbonite Availability 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 store the replica of the virtual machine from the source.



- **Current Servers**—This list contains the servers currently available in your console session. Servers that are not licensed for the workflow you have selected and those not applicable to the workload type you have selected will be filtered out of the list. Select your target server from the list. 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 full server to Hyper-V job. Configure those options that are applicable to your environment.

Go to each page identified below to see the options available for that section of the **Set Options** page. After you have configured your options, continue with the next step on page 281.

- *General* on page 264
- *Replica Virtual Machine Location* on page 264
- *Replica Virtual Machine Configuration* on page 266
- *Replica Virtual Machine Volumes* on page 267
- *Replica Virtual Machine Network Settings* on page 268
- *Failover Monitor* on page 271
- *Test Failover* on page 269
- *Failover Options* on page 272
- *Failover Identity* on page 273
- *Mirror, Verify & Orphaned Files* on page 274
- *Network Route* on page 277
- *Snapshots* on page 278
- *Compression* on page 279
- *Bandwidth* on page 280

General

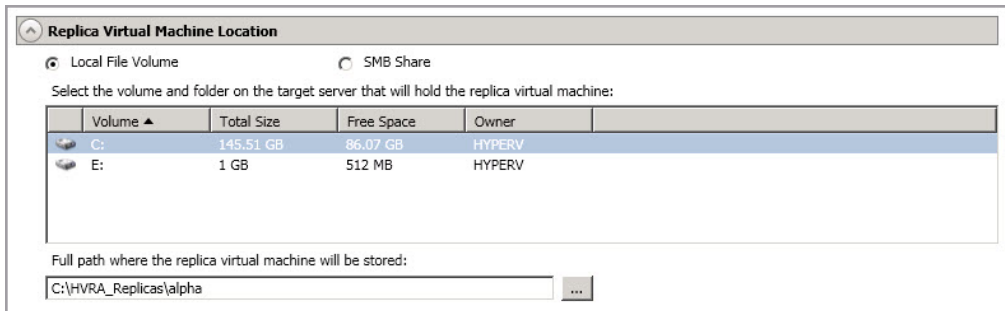


General

Job name:

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

Replica Virtual Machine Location



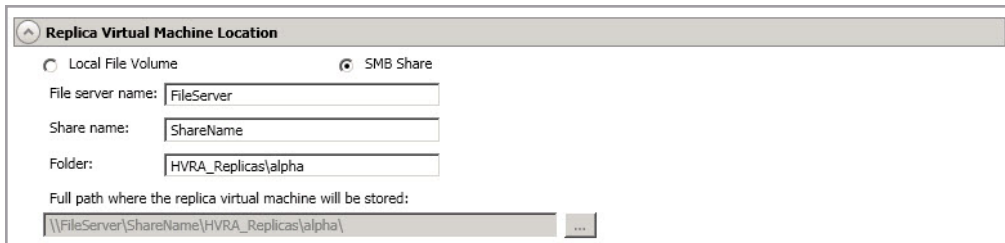
Replica Virtual Machine Location

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:	145.51 GB	86.07 GB	HYPERV
E:	1 GB	512 MB	HYPERV

Full path where the replica virtual machine will be stored:



Replica Virtual Machine Location

Local File Volume SMB Share

File server name:

Share name:

Folder:

Full path where the replica virtual machine will be stored:



The options on this page will vary depending on the operating system of your target Hyper-V server. Also, the SMB share option is only available if your target Hyper-V server is a standalone 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 Availability 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 Availability 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 Availability 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**—Depending on the operating system of your source and target Hyper-V server, you may have the choice to specify if the replica virtual machine is generation 1 or 2. If your operating system combination does not support different generations, you will not see this option.
- **Network adapter type**—Depending on the operating system of your source and target Hyper-V server, you may be able to select the type of adapter, **Legacy** or **Synthetic**, to use on the replica virtual machine. If your operating system combination does not support different adapter types, you will not see the option. This selection will apply to all adapters on the replica, and there is a limit of four legacy adapters and eight synthetic adapters.
- **Virtual switches**—Identify how you want to handle the network mapping after failover. The **Source Network Adapter** column lists the NICs from the source. Map each one to a

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 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.

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.

Snapshots are stored on the replica, so if you enable snapshots, be sure that you configure your replica virtual machine disk size large enough to maintain snapshots.

- **Replica Disk Format**—For each volume you are protecting, specify the format of the disk, **Dynamic** or **Fixed**, that will be created on the replica virtual machine. Any disk format specification will be discarded if you are reusing a disk from the **Full path where the replica virtual machine will be stored** from the **Replica Virtual Machine Location** section.
- **Storage Controller**—For each volume you are protecting, specify the type of **Storage Controller** that you want to use for each volume on the target, 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 Availability 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 Availability 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

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.



IPv6 is not supported for WAN failover.

- **Network adapters**—Select a network adapter from the source and specify the **Replica IP addresses**, **Replica Default Gateways**, and **Replica DNS Server addresses** to be used after failover. If you add multiple gateways or DNS servers, you can sort them by using the arrow up and arrow down buttons. Repeat this step for each network adapter on the source.



Updates made during failover will be based on the network adapter name when protection is established. If you change that name, you will need to delete the job and re-create it so the new name will be used during failover.

If you update one of the advanced settings (IP address, gateway, or DNS server), then you must update all of them. Otherwise, the remaining items will be left blank. If you do not specify any of the advanced settings, the replica virtual machine will be assigned the same network configuration as the source.

By default, the source IP address will be included in the target IP address list as the default address. If you do not want the source IP address to be the default address on the target after failover, remove that address from the **Replica IP addresses** list.

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.
- Scheduled snapshots will be deferred until the test replica is online.
- 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, including your test failover snapshot. (Performing a live failover or failing over to a snapshot will automatically undo any test in progress.)
- You can delete the test failover snapshot, if desired, using the **Manage Snapshots** option on the **Jobs** page.

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 failover. 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 Monitor

Failover Monitor

Total time to failure: 00:05:00

Consecutive failures: 20

Monitor on this interval: 00:00:10

Network monitoring

Monitor these addresses:

Source IP Address
<input checked="" type="checkbox"/> 172.31.206.201

Monitoring method: Network service

Failover trigger: All monitored IP addresses fail

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

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

- **Consecutive failures**—Specify how many attempts the target will make to contact the source before the source is considered failed. For example, if you have this option set to 20, and your source fails to respond to the target 20 times in a row, this will be considered a failure.
- **Monitor on this interval**—Specify, in hours:minutes:seconds, how long to wait between attempts to contact the source to confirm it is online. This means that after a response (success or failure) is received from the source, Carbonite Availability will wait the specified interval time before contacting the source again. If you set the interval to 00:00:00, then a new check will be initiated immediately after the response is received.

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

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

- **Network monitoring**—With this option, the target will monitor the source using a network ping.
 - **Monitor these addresses**—Select each **Source IP Address** that you want the target to monitor. If you are using a NAT environment, do not select a private IP address on the source because the target cannot reach the source's private address, thus causing an immediate failure. Also for NAT environments, you will see an additional field for the **Replication Service port**. This gives you the ability to specify the port number to be used with the address, allowing the target to monitor the source through a router.
 - **Monitoring method**—At this time, **Network service** is the only option. Failover monitoring will test source availability by using an ICMP ping to confirm the route is active.
 - **Failover trigger**—If you are monitoring multiple IP addresses, specify when you want a failover condition to be triggered.
 - **One monitored IP address fails**—A failover condition will be triggered when any one of the monitored IP addresses fails. If each IP address is on a different subnet, you may want to trigger failover after one fails.
 - **All monitored IP addresses fail**—A failover condition will be triggered when all monitored IP addresses fail. If there are multiple, redundant paths to a server, losing one probably means an isolated network problem and you should wait for all IP addresses to fail.

Failover Options



- **Wait for user to initiate failover**—The failover process can wait for you to initiate it, allowing you to control when failover occurs. When a failure occurs, the job will wait in **Failover Condition Met** for you to manually initiate the failover process. Disable this option if you want failover to occur immediately when a failure occurs.

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

- **Retain target network configuration**—Because the network configuration is set in the **Replica Virtual Machine Network Settings** section, the network configuration is automatically set to retain the target configuration, which is in essence retaining the replica identity. This section is essentially for updating DNS, if desired.
 - **Update DNS server**—Specify if you want Carbonite Availability to update your DNS server on failover. If DNS updates are made, the DNS records will be locked during failover. Be sure and review the job requirements for updating DNS.



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

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

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.

Mirror, Verify & Orphaned Files

- **Mirror Options**—Choose a comparison method and whether to mirror the entire file or only the bytes that differ in each file.
 - **Do not compare files. Send the entire file.**—Carbonite Availability 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 the mirror time 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 Availability 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 speed. 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 Availability 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 speed.
- **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 Availability 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 Availability 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 Availability 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 63 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 Availability 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 Availability log, and during difference mirrors, verifications, and restorations, all orphaned file processing messages are logged to the Carbonite Availability log. The orphaned file log is located in the **Logging folder** specified for the source. See *Log file properties* on page 68 for details on the location of that folder. The orphaned log file is appended to during each orphaned file processing during a mirror, and the log file will be a maximum of 50 MB.

Network Route





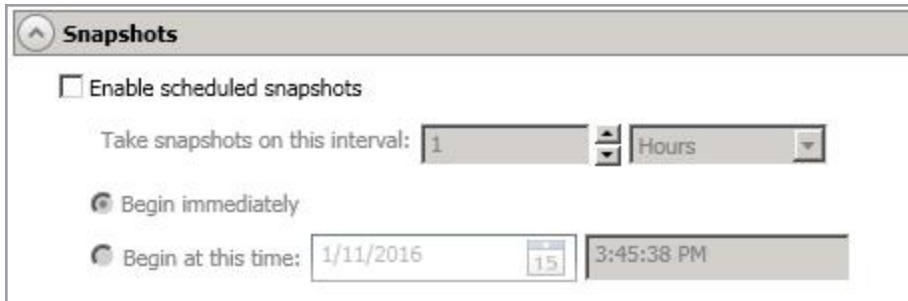
This section is not applicable to clustered environments.

By default, Carbonite Availability will select an IP address on the target 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 Availability traffic. For example, you can separate regular network traffic and Carbonite Availability 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.

Snapshots



Snapshots

Enable scheduled snapshots

Take snapshots on this interval: 1 Hours

Begin immediately

Begin at this time: 1/11/2016 3:45:38 PM

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

Turn on **Enable scheduled snapshots** if you want Carbonite Availability to take snapshots automatically at set intervals.


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



See *Managing snapshots* on page 76 for details on taking manual snapshots and deleting snapshots.

Snapshots are stored inside the replica virtual disk, so be sure that you configure your replica virtual machine disks large enough to maintain snapshots. Also, you may want to set the size limit on how much space snapshots can use. See your VSS documentation for more details on setting a size limit.

Compression



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

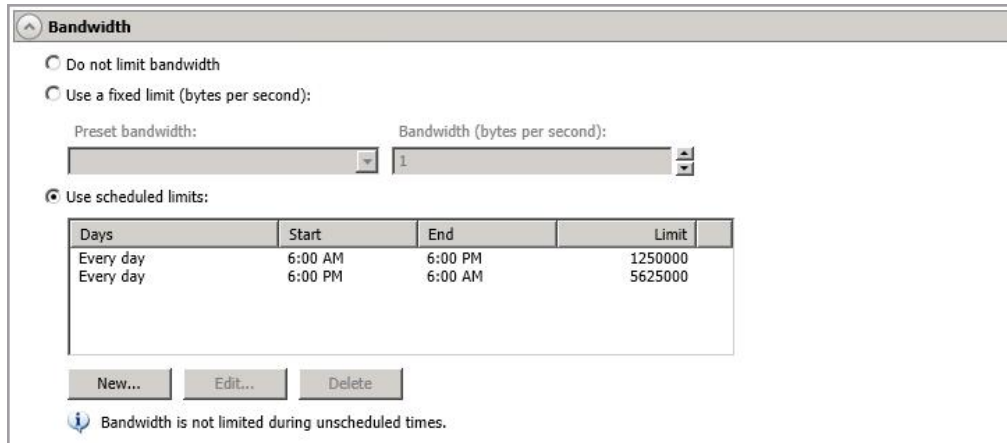
Keep in mind that the process of compressing data impacts processor usage on the source. If you notice an impact on performance while compression is enabled in your environment, either adjust to a lower level of compression, or leave compression disabled. Use the following guidelines to determine whether 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 Availability 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 Availability data transmissions. When a bandwidth limit is specified, Carbonite Availability never exceeds that allotted amount. The bandwidth not in use by Carbonite Availability is available for all other network traffic.



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

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

- **Do not limit bandwidth**—Carbonite Availability will transmit data using 100% bandwidth availability.
- **Use a fixed limit**—Carbonite Availability 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 Availability 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 Availability 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 Availability correct the problem for you. For those errors that Carbonite Availability 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 Availability log files.

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



Jobs in a NAT environment may take longer to start.

Once a job is created, do not change the name of underlying hardware components used in the job. For example, volume names, network adapter names, or 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 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 33), 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 283
- *Detailed job information displayed in the bottom right pane* on page 286
- *Job controls* on page 288


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 Availability 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 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 Availability 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.
- **Snapshot Reverted**—The data on the source and target do not match because a snapshot has been applied on the target. Restore the data from the target back to the source. If you want to discard the changes on the target, you can remirror to resynchronize the source and target.
- **Busy**—The source is low on memory causing a delay in getting the state of the data on the target.
- **Not Loaded**—Carbonite Availability 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 54.

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 54.

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. This is the current time where the console is running.

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 **Edit Job Properties** page.

Delete

Stops (if running) and deletes the selected jobs.

If you no longer want to protect the source and no longer need the replica of the source on the target, select to delete the associated replica virtual machine. Selecting this option will remove the job and completely delete the replica virtual machine on the target. Do not select this option if you want to keep the replica of the source on the target. If you do not select the delete option, the source replica will be preserved on the target.

Provide Credentials

Changes the login credentials that the job (which is on the target machine) uses to authenticate to the servers in the job. This button opens the Provide Credentials dialog box where you can specify the new account information and which servers you want to update.

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 Availability queue was not exhausted during the time the job was paused. If the Carbonite Availability queue was exhausted during the time the job was paused, the job will restart mirroring and replication.

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

Pause

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

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

Stop

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

Take Snapshot

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

Manage Snapshots

Allows you to manage your snapshots by taking and deleting snapshots for the selected job. See *Managing snapshots* on page 76 for more information.

Failover or Cutover

Starts the failover process. See *Failing over full server to Hyper-V jobs* on page 302 for the process and details of failing over a full server to Hyper-V job.

Failback

Starts the failback process. Failback does not apply to full server to Hyper-V jobs.

Restore

Starts the restoration process. Restore does not apply to full server to Hyper-V jobs.

Reverse

Reverses protection. Reverse protection does not apply to full server to Hyper-V jobs.

Undo Failover or Cutover

Cancels a test failover by undoing it. This resets the servers and the job back to their original state. See *Failing over full server to Hyper-V jobs* on page 302 for the process and details of undoing a failed over full server to Hyper-V job.

View Job Log

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

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 Availability 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 Availability 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 Availability 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 the mirror time 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 Availability 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 speed. 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 Availability 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 speed.

- **Calculate size of protected data before mirroring**—Specify if you want Carbonite Availability 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 Availability will transmit data using 100% bandwidth availability.
 - **Fixed bandwidth limit**—Carbonite Availability 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 Availability 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 Availability processing, not the entire server.

While the target is paused, the Carbonite Availability 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 Availability 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.

- **Refresh Status**—Refreshes the job status immediately.

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 33.

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 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 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.
- **Snapshot Reverted**—The data on the source and target do not match because a snapshot has been applied on the target. Restore the data from the target back to the source. If you want to discard the changes on the target, you can remirror to resynchronize the source and target.
- **Busy**—The source is low on memory causing a delay in getting the state of the data on the target.
- **Not Loaded**—Carbonite Availability 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 Availability 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 54.

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 54.

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 Availability 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 job

Over time, you may want to confirm that any changes in your network or environment have not impacted your Carbonite Availability 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 Availability 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 Availability correct the problem for you. For those errors that Carbonite Availability 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 job

Use these instructions to edit a full server to Hyper-V 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 Availability monitor security access.)
2. You will see the same options available for your full server to Hyper-V job as when you created the job, but you will not be able to edit all of them. If desired, edit those options that are configurable for an existing job. See *Creating a full server to Hyper-V job* on page 261 for details on each job option.



Changing some options may require Carbonite Availability 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 Availability 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 Availability orphan file detection cannot remove the data for you. Therefore, you have to remove it manually.

4. Click **Next** to continue.
5. Carbonite Availability 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 Availability correct the problem for you. For those errors that Carbonite Availability 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 Availability log files.

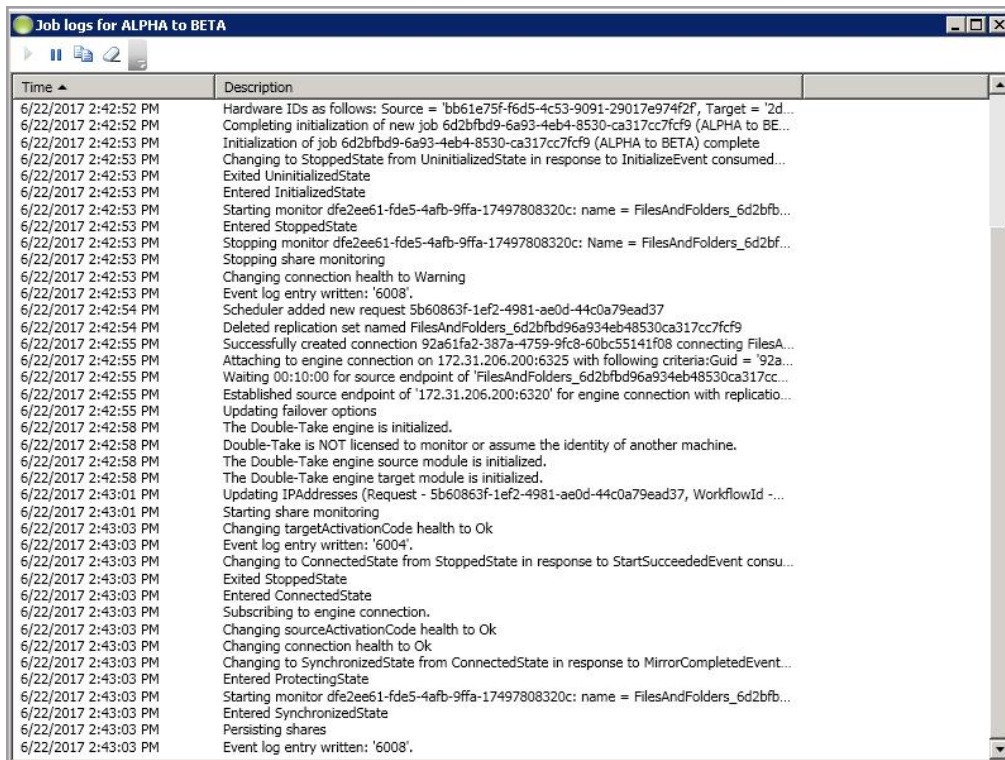
6. Once your servers have passed validation and you are ready to update your job, click **Finish**.

Viewing a full server to Hyper-V 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.



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



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



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



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

Copy

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

Clear

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

Time

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

Description

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

Failing over full server to Hyper-V jobs

When a failover condition has been met, failover will be triggered automatically if you disabled the wait for user option during your failover configuration. If the wait for user before failover option is enabled, you will be notified in the console when a failover condition has been met. At that time, you will need to trigger it manually from the console when you are ready.



If you have paused your target, failover will not start if configured for automatic failover, and it cannot be initiated if configured for manual intervention. You must resume the target before failover will automatically start or before you can manually start it.

1. On the **Jobs** page, highlight the job that you want to failover and click **Failover or Cutover** in the toolbar.
2. Select the type of failover to perform.
 - **Failover to live data**—Select this option to initiate a full, live failover using the current data on the target. This option will shutdown the source machine (if it is online), stop the protection job, and start the replica virtual machine on the target with full network connectivity.
 - **Perform test failover**—Select this option to perform a test failover.
 - 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.
 - Scheduled snapshots will be deferred until the test replica is online.
 - 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, including your test failover snapshot. (Performing a live failover or failing over to a snapshot will automatically undo any test in progress.)
 - You can delete the test failover snapshot, if desired, using the **Manage Snapshots** option on the **Jobs** page.
 - **Failover to a snapshot**—Select this option to initiate a full, live failover without using the current data on the target. Instead, select a snapshot and the data on the target will be reverted to that snapshot. This option will not be available if there are no snapshots on the target. This option is also not applicable to clustered environments. To help you understand

what snapshots are available, the **Type** indicates the kind of snapshot.

- **Scheduled**—This snapshot was taken as part of a periodic snapshot.
 - **Deferred**—This snapshot was taken as part of a periodic snapshot, although it did not occur at the specified interval because the job between the source and target was not in a good state.
 - **Manual**—This snapshot was taken manually by a user.
3. Select how you want to handle the data in the target queue.
- **Apply data in target queues before failover or cutover**—All of the data in the target queue will be applied before failover begins. The advantage to this option is that all of the data that the target has received will be applied before failover begins. The disadvantage to this option is depending on the amount of data in queue, the amount of time to apply all of the data could be lengthy.
 - **Discard data in the target queues and failover or cutover immediately**—All of the data in the target queue will be discarded and failover will begin immediately. The advantage to this option is that failover will occur immediately. The disadvantage is that any data in the target queue will be lost.
 - **Revert to last good snapshot if target data state is bad**—If the target data is in a bad state, Carbonite Availability will automatically revert to the last good Carbonite Availability snapshot before failover begins. If the target data is in a good state, Carbonite Availability will not revert the target data. Instead, Carbonite Availability will apply the data in the target queue and then failover. The advantage to this option is that good data on the target is guaranteed to be used. The disadvantage is that if the target data state is bad, you will lose any data between the last good snapshot and the failure.
4. When you are ready to begin failover, click **Failover**.



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

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

- **Retail licensing**—Retail licensing allows the activation of a single operating system installation.
 1. Open the **System** applet in Windows **Control Panel**.
 2. Under **Windows activation** at the bottom of the page, click **Change product key**.
 3. Enter your retail license key. You may need access to the Internet or to call Microsoft to complete the activation.
- **MAK volume licensing**—Multiple Activation Key (MAK) licensing allows the activation of multiple operating system installations using the same activation key.

1. View or download the Microsoft Volume Activation Deployment Guide from the Microsoft web site.
 2. Using an administrative user account, open a command prompt and follow the instructions from the deployment guide to activate MAK clients. Multiple reboots may be necessary before you can access a command prompt. You may need access to the Internet or to call Microsoft to complete the activation.
- **KMS volume licensing**—Key Management Service (KMS) licensing allows IT professionals to complete activations on their local network without contacting Microsoft.
 1. View or download the Microsoft Volume Activation Deployment Guide from the Microsoft web site.
 2. Using an administrative user account, open a command prompt and follow the instructions from the deployment guide to convert a MAK activation client to a KMS client. Multiple reboots may be necessary before you can access a command prompt.

Depending on your replica configuration, you may have to reboot your replica after failover. You will be prompted to reboot if it is necessary.

After failover, if you attempt to use a full server job to revert back to your original configuration, you will need to perform a few additional tasks before creating the full server job. Contact technical support if you need assistance with these steps.

1. On either the source or target, stop the Double-Take and Double-Take Management Service services.
2. Remove the GUID value from HKEY_LOCAL_MACHINE\SOFTWARE\NSI Software\Double-Take\CurrentVersion\Communication\Uniqueld. Do not delete the Uniqueld key. Only delete the GUI value within the key.
3. Restart the the Double-Take and Double-Take Management Service services.
4. Remove and then add your servers back into the Carbonite Replication Console.
5. Install a different license on the original source and complete a host transfer if necessary.

If your source is running on Windows Server 2008 and the target replica has one or more SCSI drives, then after failover the CD/DVD ROM will not be allocated. If the CD/DVD ROM is required, you will need to edit the virtual machine settings to add a CD/DVD ROM after failover. By not allocating a CD/DVD ROM under these specific conditions, drive letter consistency will be guaranteed.

If your source is a virtual machine running on an ESX host and you have VMware Tools installed, VMware Tools will be disabled during failover.

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, one or more 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 or remove these hidden NICs.

5. If you performed a test failover, you can undo it by selecting **Undo Failover or Cutover** in the toolbar. Confirm the undo process when prompted. The replica virtual machine on the target will be shut down and, if configured, the virtual disks used for the test failover will be deleted.
6. There is no reverse or failback once you have completed a live failover. If you need to go back to your original hardware, delete the job and re-create a new one if your original source was a virtual server. If your original source was a physical server, you will need to use a full server job.

Reversing protection after failover for full server to Hyper-V jobs

There is no automated reverse or failback for a full server to Hyper-V job once you have failed over. If you need to go back to your original hardware, you will need to create a new job in the opposite direction following one of the processes below, depending on the original source.

- **Physical server**—Use these steps if your original source is a physical server.
 1. Resolve the problems on the original source that caused it to fail. If you need to deploy a new server, use the same operating system and disk configuration as the original source.
 2. If Carbonite Availability is still running on the original source, replace the license since that license is currently running on the failed over server. If Carbonite Availability is no longer installed, reinstall it with an appropriate license.
 3. Create a full server job from the failed over server to your original source. See *Creating a full server job* on page 149 for details on creating this job.
 4. Once the initial mirror is complete, failover the full server job. See *Failing over full server jobs* on page 192 for details on this process.
- **Hyper-V virtual server**—Use these steps if your original source is a virtual server on a Hyper-V host.
 1. Delete the original source virtual server from the Hyper-V host. If you want to reuse the .vhd files again, only delete the original source virtual server from the Hyper-V inventory.
 2. If it is not already, install and license Carbonite Availability on the Hyper-V host where the original source virtual server is located. The host will be the target of the new job you are going to create.
 3. Create a full server to Hyper-V job from your failed over server to the Hyper-V host where the original source virtual server is located. See *Creating a full server to Hyper-V job* on page 261 for details on creating this job.
 4. Once the initial mirror is complete, failover the full server to Hyper-V job. See *Failing over full server to Hyper-V jobs* on page 302 for details on this process.
- **ESX virtual server**—Use these steps if your original source is a virtual server on an ESX host.
 1. Delete the original source virtual server from the ESX host. If you want to reuse the .vmdk files again, only delete the original source virtual server from the ESX inventory.
 2. If you do not have one already, create a virtual recovery appliance on the ESX host where the original source virtual server is located. This appliance will be the target of the new job you are going to create. This appliance needs Carbonite Availability installed and licensed on it. For more details, see *Full server to ESX requirements* on page 308.
 3. Create a full server to ESX job from your failed over server to the appliance on the ESX host where the original source virtual server is located. See *Creating a full server to ESX job* on page 315 for details on creating this job.
 4. Once the initial mirror is complete, failover the full server to ESX job. See *Failing over full server to ESX jobs* on page 360 for details on this process.

Chapter 9 Full server to ESX protection

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

- *Full server to ESX requirements* on page 308—Full server to ESX protection includes specific requirements for this type of protection.
- *Creating a full server to ESX job* on page 315—This section includes step-by-step instructions for creating a full server to ESX job.
- *Managing and controlling full server to ESX jobs* on page 341—You can view status information about your full server to ESX jobs and learn how to control these jobs.
- *Failing over full server to ESX jobs* on page 360—Use this section when a failover condition has been met or if you want to failover manually.
- *Reversing protection after failover for full server to ESX jobs* on page 364—Use this section to create a reverse job. A reverse job is used after a live or snapshot failover. It takes the replica virtual machine you failed over to and creates a new protection job from that replica virtual machine to a new replica virtual machine on an ESX host.

Full server to ESX requirements

Use these requirements for full server to ESX protection.

- **Source server**—The following operating systems are supported for on the source for full server to ESX jobs.
 - Windows 2019 and Server Core 2019
 - 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 and 2019 support are for the primary operating system features available in Windows 2012. New operating system features specific to Windows 2016 and 2019, such as Nano Server, Windows Containers, and so on, are not supported.

Due to VMware limitations, Windows 2016 and 2019 are not supported on ESX 5.5.

DNS updates are not supported for Server Core servers.

If your source is a Hyper-V server, you will be able to protect it, however the Hyper-V role and features will not be available after failover.

-
- **Target host server**—The target host server must be an ESX server. It can be any of the following ESX operating systems.
 - ESXi 6.5
 - ESXi 6.7
 - ESXi 7.0



The free versions of ESX restrict functionality that Carbonite Availability 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 and Storage vMotion are supported with the following caveats. DRS (Distributed Resource Scheduler) is not supported.
 - **Source**—The source supports both Host and Storage vMotion as long as you are using vCenter.
 - **Appliance**—The appliance supports both Host and Storage vMotion as long as you are using vCenter, except the storage location should not be moved once the failover or test failover process has started. If the storage location is moved during failover or test failover, the process could potentially fail.

- **Replica**—The replica supports both Host and Storage vMotion as long as you are using vCenter and the failover or test failover process has completed.
- **Virtual recovery appliance**—The ESX server must have an existing virtual machine, known as a virtual recovery appliance, that meets the following requirements. (When you establish protection, the virtual recovery appliance will create a new virtual server, mount disks, format disks, and so on. If failover occurs, the new virtual machine is detached from the virtual recovery appliance and powered on. Once the new virtual machine is online, it will have the identity, data, and system state of the source. Since the virtual recovery appliance maintains its own identity, it can be reused for additional failovers.)
 - **Operating system**—The virtual recovery appliance can be any of the operating systems listed above for the source server in the following combinations.
 - **2019, 2016, or 2012 R2**—If your appliance is Windows 2019, 2016 or 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 and later is not supported on ESX 5.5.

If your source is Windows 2016 and you are using ReFS volumes, your virtual recovery appliance must be Windows 2016. You cannot use a Windows 2019 virtual recovery appliance because the formatting of ReFS is different between Windows 2016 and Windows 2019.

- **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 Availability**—The virtual recovery appliance must have Carbonite Availability installed and licensed on it.
- **Maximum number of protected disks**—A single virtual recovery appliance can protect a maximum of 59 disks from any number of sources.
- **Snapshots**—Do not take snapshots of the virtual recovery appliance, because they will interfere with proper failover.
- **Permissions**—If you want to limit the permissions required for the account that you will be using for your full server to ESX job, your account must have at a minimum the permissions listed below. These permissions can be set at the vCenter, Datacenter, or host level.
 - **Datastore**—Allocate Space, Browse Datastore, Low level file operations, and Remove File
 - **Host, Local Operations**—Create Virtual Machine, Delete Virtual Machine, and Reconfigure virtual machine
 - **Network**—Assign Network
 - **Resource**—Assign virtual machine to resource pool
 - **Scheduled Task**—Create Tasks, Modify Task, Remove Task, and Run Task
 - **Tasks**—Create task and Update task
 - **Virtual Machine, Configuration**—Add existing disk, Add new disk, Add or remove device, Change resource, Modify device settings, and Remove disk

- **Virtual Machine, Interaction**—Device connection, Power off, and Power on
- **Virtual Machine, Inventory**—Create new, Register, Remove, and Unregister

Make sure if you also define permissions at the VMs and Templates level in vCenter that you have not denied any of the required permissions listed above.

- **Domain controllers**—If your source is a domain controller, it will start in a non-authoritative restore mode after failover. This means that if the source was communicating with other domain controllers before failover, it will require one of those domain controllers to be reachable after failover so it can request updates. If this communication is not available, the domain controller will not function after failover. If the source is the only domain controller, this is not an issue.
- **File system**—Carbonite Availability supports the NTFS file system. On Windows 2016 and later, 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 21.



Because ReFS is formatted differently on Windows 2016 and Windows 2019, you cannot use a Windows 2016 source with ReFS to a Windows 2019 target.

-
- **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 Availability program files.

-
- **Server name**—Carbonite Availability includes Unicode file system support, but your server name must still be in ASCII format. Additionally, all Carbonite Availability servers and appliances must have a unique server name.



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

-
- **Time**—The clock on your Carbonite Availability servers must be within a few minutes of each other, relative to UTC. Large time skews (more than five minutes) will cause Carbonite Availability 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.
 - WAN failover is not supported with IPv6 addresses.

- If you are using IPv6 on your servers, your console 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 Availability 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 Availability.
- **Network adapters**—Your source can have no more than ten NICs enabled.
- **NAT support**—Carbonite Availability 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 Availability 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 Availability 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 Availability installation program will automatically attempt to configure ports 6320, 6325, and 6326 for Carbonite Availability. 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 367 for instructions on handling firewalls in your environment.

- **Windows Management Instrumentation (WMI)**—Carbonite Availability is dependent on the WMI service. If you do not use this service in your environment, contact technical support.
- **Snapshots**—You can take and failover to Carbonite Availability snapshots using a full server to ESX job.

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

- **Snapshot location**—Snapshots are taken and stored inside the replica virtual disk, so be sure that you configure your replica virtual machine disks large enough to maintain snapshots.
- **Carbonite Availability installation location**—In order to enable Carbonite Availability snapshots, Carbonite Availability must be installed on the system drive. If Carbonite Availability is not installed on the system drive, snapshots will be disabled when enabling protection.
- **Server IP address**—If you have specified an IP address as the source server name, but that IP address is not the server's primary IP address, you will have issues with snapshot functionality. If you need to use snapshots, use the source's primary IP address or its name.
- **Snapshot limitations**—Sometimes taking a snapshot may not be possible. For example, there may not be enough disk space to create and store the snapshot, or maybe the target is too low on memory. If a snapshot fails, an Event message and a Carbonite Availability log message are both created and logged.

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

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

You must also keep in mind that if you are using extended functionality provided by Volume Shadow Copy, you need to be aware of the impacts that functionality may have on Carbonite Availability. For example, if you change the location where the shadow copies are stored and an error occurs, it may appear to be a Carbonite Availability error when it is

in fact a Volume Shadow Copy error. Be sure and review any events created by the VolSnap driver and check your Volume Shadow Copy documentation for details.

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

Volume Shadow Copy snapshots are associated with the volume they belong to. Since Carbonite Availability 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.

- **Supported configurations**—The following table identifies the supported configurations for a full server to ESX job.

Server to Host Configuration	Description	Supported	Not Supported
One to one active/standby	You can protect a single source to a single target host.	X	
One to one active/active	This configuration (where both the source and target use the same job type to actively replicate to each other) is not supported and not applicable because the target is a hypervisor host.		X
Many to one	You can protect 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.	X	
One to many	You can protect a single source to multiple target hosts. The source is the only server actively replicating data. This will create redundant copies of your source.	X	
Chained	This configuration (where the source replicates to the target and then the target uses the same job type to replicate the source to a final target) is not supported and not applicable because the middle target is a hypervisor host.		X
Single server	You cannot protect a single source to itself.		X

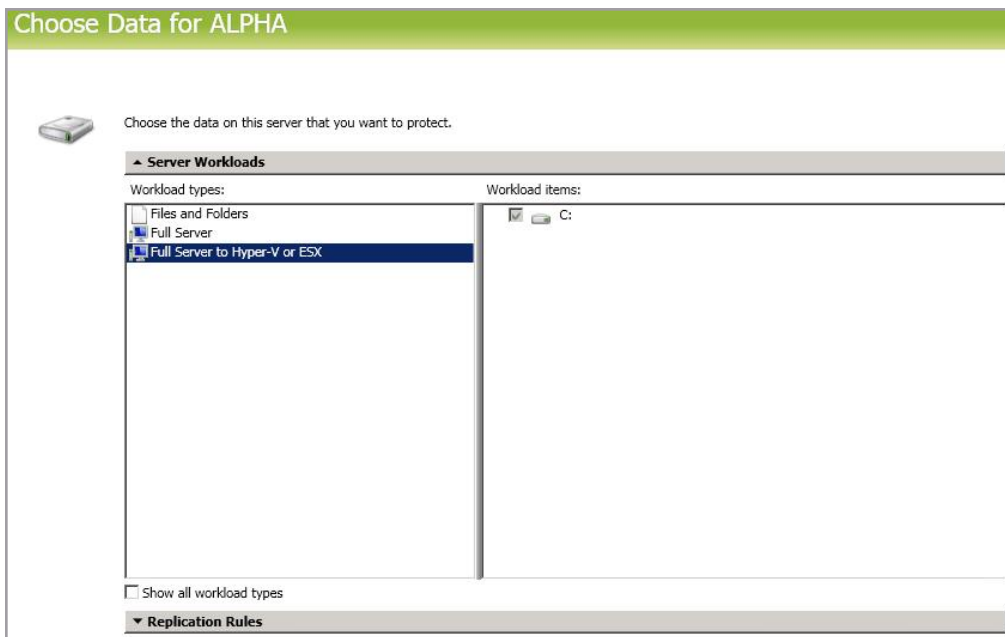
Server to Host Configuration	Description	Supported	Not Supported
Standalone to standalone	Your source and target host can be in a standalone to standalone configuration.	X	
Standalone to cluster	Your source and target host cannot be in a standalone to cluster configuration.		X
Cluster to standalone	Your source and target host cannot be in a cluster to standalone configuration.		X
Cluster to cluster	Your source and target host cannot be in a cluster to cluster configuration.		X

Creating a full server to ESX job

Use these instructions to create a full server to ESX job.

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

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



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

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

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


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

If you need to remove a rule, highlight it in the list at the bottom and click **Remove Rule**. Be careful when removing rules. Carbonite Availability 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 requirements* on page 308 for details on the virtual recovery appliance.



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

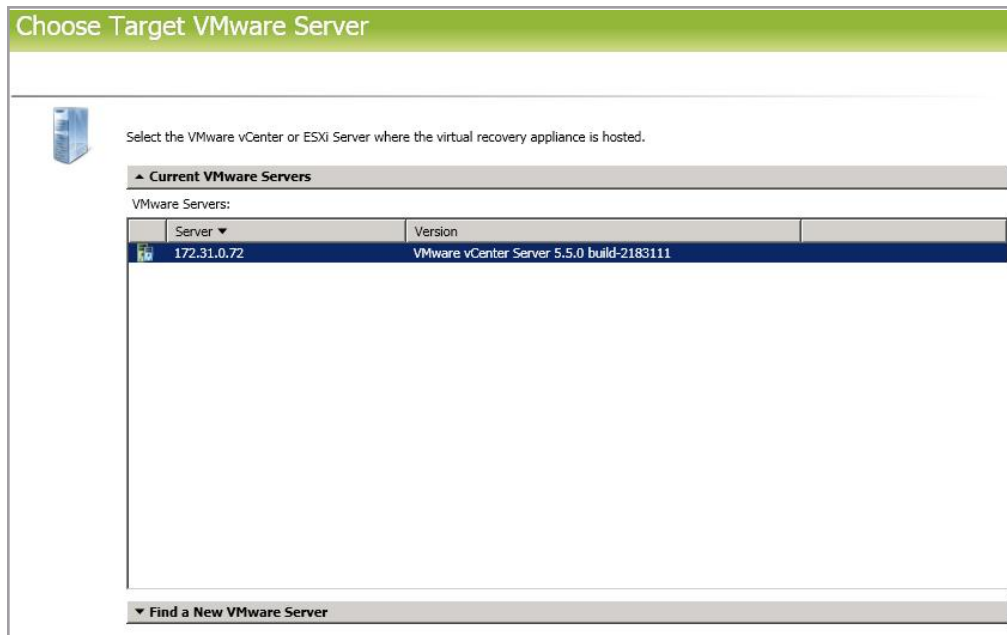
- **Find a New Server**—If the server you need is not in the **Current Servers** list, click the **Find a New Server** heading. From here, you can specify a server along with credentials for logging in to the server. If necessary, you can click **Browse** to select a server from a network drill-down 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 full server to ESX job. Configure those options that are applicable to your environment.

Go to each page identified below to see the options available for that section of the **Set Options** page. After you have configured your options, continue with the next step on page 339.

- *General* on page 318
- *Replica Virtual Machine Location* on page 319
- *Replica Virtual Machine Configuration* on page 319
- *Replica Virtual Machine Volumes* on page 320
- *Replica Virtual Machine Network Settings* on page 322
- *Failover Monitor* on page 325
- *Test Failover* on page 323
- *Failover Options* on page 326
- *Failover Identity* on page 327
- *Mirror, Verify & Orphaned Files* on page 328
- *Network Route* on page 331
- *Snapshots* on page 332
- *Compression* on page 333
- *Bandwidth* on page 334

General



The screenshot shows a dialog box titled "General" with a close button (X) in the top-left corner. Below the title bar, there is a label "Job name:" followed by a text input field containing the text "alpha to beta".

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

Replica Virtual Machine Location

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

Display name:
Alpha_Replica

Hardware configuration:

	Source	Replica
Sockets	2	2
Cores per socket		1
Memory (MB)	4096	4096

Replica boot system:
BIOS

Network adapter type:
E1000

Virtual switches:

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

- **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.

- **Replica boot system**—Select the type of boot system to use on the replica virtual machine. If your source disk is larger than 2 TB, you need to select EFI.
- **Network adapter type**—Depending on the operating system of your source, you may be able to select the type of adapter to use on the replica virtual machine. This selection will apply to all adapters on the replica.



If the operating system on the source is not compatible with the VmxNet3 driver on the target appliance, and the source does not have VMware Tools already, you need to install VMware Tools on the replica after failover in order for the VmxNet3 adapter to work correctly. Alternatively, you could select a different network adapter type, if another type is available.

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

Replica Virtual Machine Volumes

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. Any disk size specification will be discarded if you will be using an existing disk. .



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.

Snapshots are stored on the replica, so if you enable snapshots, be sure that you configure your replica virtual machine disk size large enough to maintain snapshots.

- **Replica Disk Format**—For each volume you are protecting, specify the format of the disk that will be created. Any disk format specification will be discarded if you will be using an

existing disk.

- **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 Availability 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 Availability 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 Availability 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 Availability 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

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.



IPv6 is not supported for WAN failover.

If you use advanced settings, you should not configure reverse protection on the next page of the job creation wizard. Instead of the reverse protection, you will have to create a new job in order to reverse protection.

- **Network adapters**—Select a network adapter from the source and specify the **Replica IP addresses**, **Replica Default Gateways**, and **Replica DNS Server addresses** to be used after failover. If you add multiple gateways or DNS servers, you can sort them by using the arrow up and arrow down buttons. Repeat this step for each network adapter on the source.



Updates made during failover will be based on the network adapter name when protection is established. If you change that name, you will need to delete the job and re-create it so the new name will be used during failover.

If you update one of the advanced settings (IP address, gateway, or DNS server), then you must update all of them. Otherwise, the remaining items will be left blank. If you do not specify any of the advanced settings, the replica virtual machine will be assigned the same network configuration as the source.

By default, the source IP address will be included in the target IP address list as the default address. If you do not want the source IP address to be the default address on the target after failover, remove that address from the **Replica IP addresses** list.

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.
- Scheduled snapshots will be deferred until the test replica is online.
- 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, including your test failover snapshot. (Performing a live failover or failing over to a snapshot will automatically undo any test in progress.)
- You can delete the test failover snapshot, if desired, using the **Manage Snapshots** option on the **Jobs** page.

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
Local Area Connection	VM Network 5

Configuration

Configure the volumes for the test replica virtual machine.

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

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.
 - **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 failover. 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 Monitor

Source IP Address
<input checked="" type="checkbox"/> 172.31.206.201

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

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

- **Consecutive failures**—Specify how many attempts the target will make to contact the source before the source is considered failed. For example, if you have this option set to 20, and your source fails to respond to the target 20 times in a row, this will be considered a failure.
- **Monitor on this interval**—Specify, in hours:minutes:seconds, how long to wait between attempts to contact the source to confirm it is online. This means that after a response (success or failure) is received from the source, Carbonite Availability will wait the specified interval time before contacting the source again. If you set the interval to 00:00:00, then a new check will be initiated immediately after the response is received.

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

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

- **Network monitoring**—With this option, the target will monitor the source using a network ping.
 - **Monitor these addresses**—Select each **Source IP Address** that you want the target to monitor. If you are using a NAT environment, do not select a private IP address on the source because the target cannot reach the source's private address, thus causing an immediate failure. Also for NAT environments, you will see an additional field for the **Replication Service port**. This gives you the ability to specify the port number to be used with the address, allowing the target to monitor the source through a router.
 - **Monitoring method**—At this time, **Network service** is the only option. Failover monitoring will test source availability by using an ICMP ping to confirm the route is active.
 - **Failover trigger**—If you are monitoring multiple IP addresses, specify when you want a failover condition to be triggered.
 - **One monitored IP address fails**—A failover condition will be triggered when any one of the monitored IP addresses fails. If each IP address is on a different subnet, you may want to trigger failover after one fails.
 - **All monitored IP addresses fail**—A failover condition will be triggered when all monitored IP addresses fail. If there are multiple, redundant paths to a server, losing one probably means an isolated network problem and you should wait for all IP addresses to fail.

Failover Options



- **Wait for user to initiate failover**—The failover process can wait for you to initiate it, allowing you to control when failover occurs. When a failure occurs, the job will wait in **Failover Condition Met** for you to manually initiate the failover process. Disable this option if you want failover to occur immediately when a failure occurs.

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

- **Retain target network configuration**—Because the network configuration is set in the **Replica Virtual Machine Network Settings** section, the network configuration is automatically set to retain the target configuration, which is in essence retaining the replica identity. This section is essentially for updating DNS, if desired.
 - **Update DNS server**—Specify if you want Carbonite Availability to update your DNS server on failover. If DNS updates are made, the DNS records will be locked during failover. Be sure and review the job requirements for updating DNS.



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

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

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.

Mirror, Verify & Orphaned Files

- **Mirror Options**—Choose a comparison method and whether to mirror the entire file or only the bytes that differ in each file.
 - **Do not compare files. Send the entire file.**—Carbonite Availability 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 the mirror time 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 Availability 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 speed. 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 Availability 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 speed.
- **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 Availability 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 Availability 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 Availability 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.

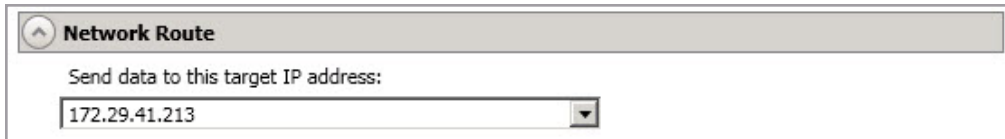
The orphaned file feature does not delete alternate data streams. To do this, use a full mirror, which will delete the additional streams when the file is re-created.

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 63 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 Availability 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 Availability log, and during difference mirrors, verifications, and restorations, all orphaned file processing messages are logged to the Carbonite Availability log. The orphaned file log is located in the **Logging folder** specified for the source. See *Log file properties* on page 68 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

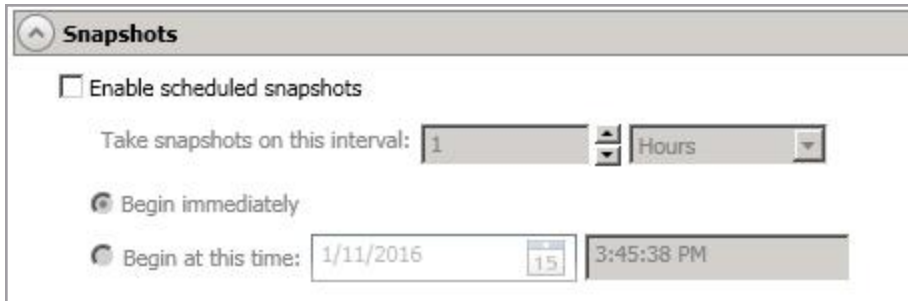


By default, Carbonite Availability will select an IP address on the target for transmissions. If desired, specify an alternate route on the target that the data will be transmitted through. This allows you to select a different route for Carbonite Availability traffic. For example, you can separate regular network traffic and Carbonite Availability 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.

Snapshots



Snapshots

Enable scheduled snapshots

Take snapshots on this interval: 1 Hours

Begin immediately

Begin at this time: 1/11/2016 3:45:38 PM

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

Turn on **Enable scheduled snapshots** if you want Carbonite Availability to take snapshots automatically at set intervals.


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



See *Managing snapshots* on page 76 for details on taking manual snapshots and deleting snapshots.

Snapshots are stored inside the replica virtual disk, so be sure that you configure your replica virtual machine disks large enough to maintain snapshots. Also, you may want to set the size limit on how much space snapshots can use. See your VSS documentation for more details on setting a size limit.

Compression



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

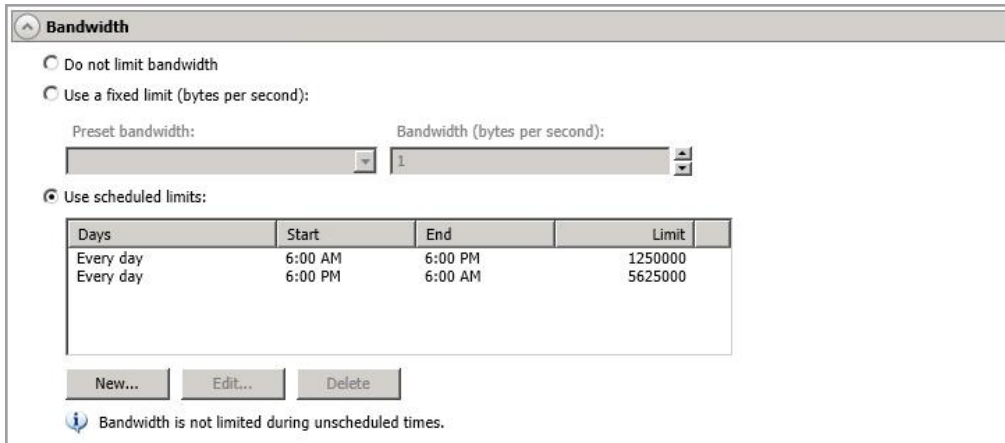
Keep in mind that the process of compressing data impacts processor usage on the source. If you notice an impact on performance while compression is enabled in your environment, either adjust to a lower level of compression, or leave compression disabled. Use the following guidelines to determine whether 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 Availability 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 Availability data transmissions. When a bandwidth limit is specified, Carbonite Availability never exceeds that allotted amount. The bandwidth not in use by Carbonite Availability is available for all other network traffic.



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

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

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



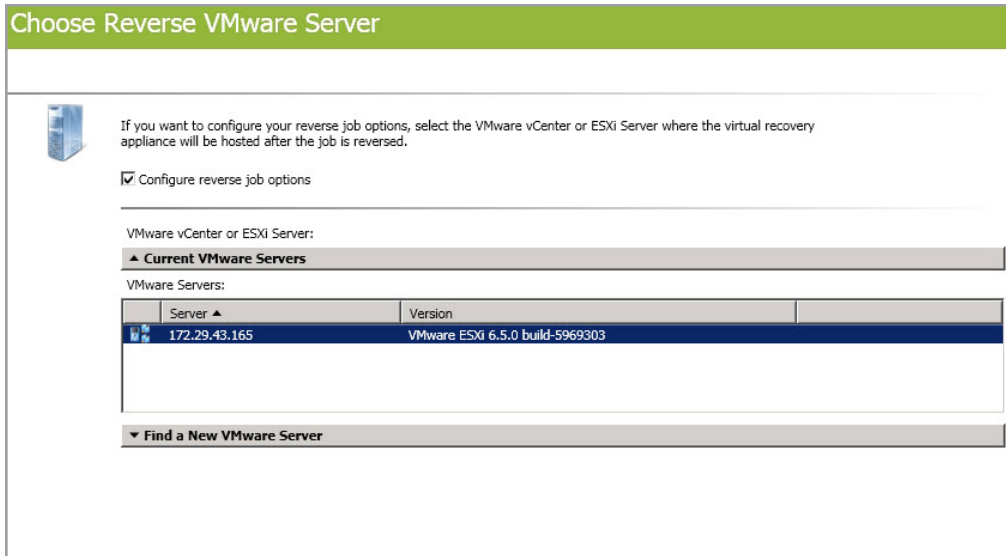
If you change your job option from **Use scheduled limits** to **Do not limit bandwidth** or **Use a fixed limit**, any schedule that you created will be preserved. That schedule will be reused if you change your job option back to **Use scheduled limits**.

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. You have the option to configure a reverse job. A reverse job is used after a live or snapshot failover. It takes the replica virtual machine you failed over to and creates a new protection job from that replica virtual machine to a new replica virtual machine on an ESX host.
-



If you selected advanced settings under **Replica Virtual Machine Network Settings** when setting your forward protection, you should not configure reverse protection. Instead of the reverse protection, you will have to create a new job in order to reverse protection.



- **Configure reverse job options**—Select this option if you want to configure the reverse job options now. You can edit any configured reverse job options later by editing the forward job. If this option is not selected, the reverse job options will be skipped at this time. You can configure reverse job options after the forward job is created by editing the job.
 - **VMware vCenter or ESXi Server** —If you are configuring the reverse job options now, select the server where the target appliance of the reverse job 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.

12. Click **Next** to continue.

13. If you are configuring reverse job options now, select the options for the reverse job. Any options you configure can be edited later by editing the forward job options.

Go to each page identified below to see the options available for that section of the **Set Reverse Job Options** page. After you have configured your options, continue with the next step on page 339.

- *Replica Virtual Machine Location* on page 337
- *Replica Virtual Machine Configuration* on page 338

- *Replica Virtual Machine Volumes* on page 338



Any job options not available for the reverse job, like mirroring options or compression, will use the default setting of a new full server to ESX job.

Replica Virtual Machine Location

Replica Virtual Machine Location

Select the virtual recovery appliance on the reverse ESX server that will be used when reversing the job:

Name ▲	Guest Host Name	Operating System	Credentials
WinAppliance	WinAppliance	Microsoft Windows Server 2012 (64-...	ldom\administr...

Credentials...

Send data to the reverse appliance using this route:

172.29.41.212

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

Volume	Total Size	Provisioned Space	Free Space	Owner
RNDESX65-DS1	399.75 GB	526.47 GB	61.09 GB	ESX65
RNDESX65-DS2	499.75 GB	849.97 GB	63.87 GB	ESX65
RNDESX65-DS3	499.75 GB	385.59 GB	114.16 GB	ESX65

- **Select the virtual recovery appliance on the reverse ESX server that will be used when reversing the job**—Select a server that will be your target appliance of the reverse job. See *Full server to ESX requirements* on page 308 for details on the virtual recovery appliance.

You will need to provide **Credentials** for the server. Specify a user that is a member of the local Double-Take Admin security group. 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.

- **Send data to the reverse appliance using this route**—By default, Carbonite Availability 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 Availability traffic. For example, you can separate regular network traffic and Carbonite Availability traffic on a machine with multiple IP addresses.



The IP address used on the replica virtual machine acting as the source will be determined through the Windows route table.

- **Select the datastore on the reverse ESX server that will hold the reverse virtual machine**—Select one of the volumes from the list to indicate the volume on the reverse

ESX server where you want to store the configuration files for the new reverse virtual server when it is created. The volume must have enough **Free Space**. You can select the location of the .vmdk files under **Replica Virtual Machine Volumes**.

Replica Virtual Machine Configuration

Reverse virtual machine display name:
Alpha_ReverseReplica

Specify the display name of the reverse virtual machine on the reverse ESX host.

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 reversing, specify the size of the replica disk on the target. Be sure and include the value in MB or GB for the disk. The value must be at least the size of the specified **Used Space** on that volume. Any disk size specification will be discarded if you will be using an existing disk.



In some cases, the reverse virtual machine may use more virtual disk space than the size of the replica disk volume due to differences in how the reverse virtual disk's block size is formatted. To avoid this issue, make sure the reverse virtual machine 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 reversing, specify the format of the disk that will be created. Any disk format specification will be discarded if you will be using an existing disk.
 - **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 reversing, specify the datastore on the where you want to store the reverse virtual disk files for the reverse 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 Availability to create a new disk for your reverse 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 Availability 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 Availability 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 Availability 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 Datastore** that you want to reuse.
14. Click **Next** to continue.
 15. Carbonite Availability 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 Availability correct the problem for you. For those errors that Carbonite Availability 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 Availability log files.

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



Jobs in a NAT environment may take longer to start.

Once a job is created, do not change the name of underlying hardware components used in the job. For example, volume names, network adapter names, or 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 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 33), 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 341
- *Detailed job information displayed in the bottom right pane* on page 344
- *Job controls* on page 346


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 Availability 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 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 Availability 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.
- **Snapshot Reverted**—The data on the source and target do not match because a snapshot has been applied on the target. Restore the data from the target back to the source. If you want to discard the changes on the target, you can remirror to resynchronize the source and target.
- **Busy**—The source is low on memory causing a delay in getting the state of the data on the target.
- **Not Loaded**—Carbonite Availability 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 54.

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 54.

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. This is the current time where the console is running.

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 **Edit Job Properties** page.

Delete

Stops (if running) and deletes the selected jobs.

If you no longer want to protect the source and no longer need the replica of the source on the target, select to delete the associated replica virtual machine. Selecting this option will remove the job and completely delete the replica virtual machine on the target. Do not select this option if you want to keep the replica of the source on the target. If you do not select the delete option, the source replica will be preserved on the target.

If you are using vCenter, but created a job directly to an ESX host, you will have an orphaned virtual machine in vCenter if you choose to delete the virtual machine. That is because the ESX host is not forwarding the delete to the vCenter. You will need to manually delete the orphaned virtual machine in vCenter.

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.

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 Availability queue was not exhausted during the time the job was paused. If the Carbonite Availability queue was exhausted during the time the job was paused, the job will restart mirroring and replication.

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

Pause

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

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

Stop

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

Take Snapshot

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

Manage Snapshots

Allows you to manage your snapshots by taking and deleting snapshots for the selected job. See *Managing snapshots* on page 76 for more information.

Failover or Cutover

Starts the failover process. See *Failing over full server to ESX jobs* on page 360 for the process and details of failing over a full server to ESX job.

Failback

Starts the failback process. Failback does not apply to full server to ESX jobs.

Restore

Starts the restoration process. Restore does not apply to full server to ESX jobs.

Reverse

Reverses protection. Reverse protection does not apply to full server to ESX jobs.

Undo Failover or Cutover

Cancels a test failover by undoing it. This resets the servers and the job back to their original state. See *Failing over full server to ESX jobs* on page 360 for the process and details of undoing a failed over full server to ESX job.

View Job Log

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

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 Availability 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 Availability 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 Availability 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 the mirror time 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 Availability 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 speed. 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 Availability 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 speed.
- **Calculate size of protected data before mirroring**—Specify if you want Carbonite Availability 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 Availability will transmit data using 100% bandwidth availability.
 - **Fixed bandwidth limit**—Carbonite Availability 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 Availability 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 Availability processing, not the entire server.

While the target is paused, the Carbonite Availability 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 Availability 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.

- **Refresh Status**—Refreshes the job status immediately.

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 33.

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 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 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.
- **Snapshot Reverted**—The data on the source and target do not match because a snapshot has been applied on the target. Restore the data from the target back to the source. If you want to discard the changes on the target, you can remirror to resynchronize the source and target.
- **Busy**—The source is low on memory causing a delay in getting the state of the data on the target.
- **Not Loaded**—Carbonite Availability 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 Availability 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 54.

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 54.

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 Availability 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 job

Over time, you may want to confirm that any changes in your network or environment have not impacted your Carbonite Availability 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 Availability 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 Availability correct the problem for you. For those errors that Carbonite Availability 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 job

Use these instructions to edit a full server to ESX 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 Availability monitor security access.)
2. If your job has not yet failed over, you will see the same options for your full server to ESX job as when you created the job, but you will not be able to edit all of them. If desired, edit those options that are configurable for an existing job. See *Creating a full server to ESX job* on page 315 for details on each job option.



Changing some options may require Carbonite Availability 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.

If your job has already failed over, you will skip the job options for the forward job and go directly to the selection for editing the reverse job options.

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 Availability 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 Availability orphan file detection cannot remove the data for you. Therefore, you have to remove it manually.

4. Click **Next** to continue.
5. You have the option to edit a previously configured reverse job, configure a reverse job that you did not previously configure, or skip the reverse job configuration. A reverse job is used after a live or snapshot failover. It takes the replica virtual machine you failed over to and creates a new protection job from that replica virtual machine to a new replica virtual machine on an ESX host. If you want to configure or edit the reverse job, make sure the configure option is selected and click **Next**. If you do not want to configure reverse job options, do not select the option and click **Next**.
6. If you selected to configure reverse options, you will be presented with the reverse settings. These are the same options as when you created the job. See *Creating a full server to ESX job* on page 315 for details on each reverse job option.
7. Click **Next** to continue.
8. Carbonite Availability 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 Availability correct the problem for you. For those errors that Carbonite Availability 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 Availability log files.

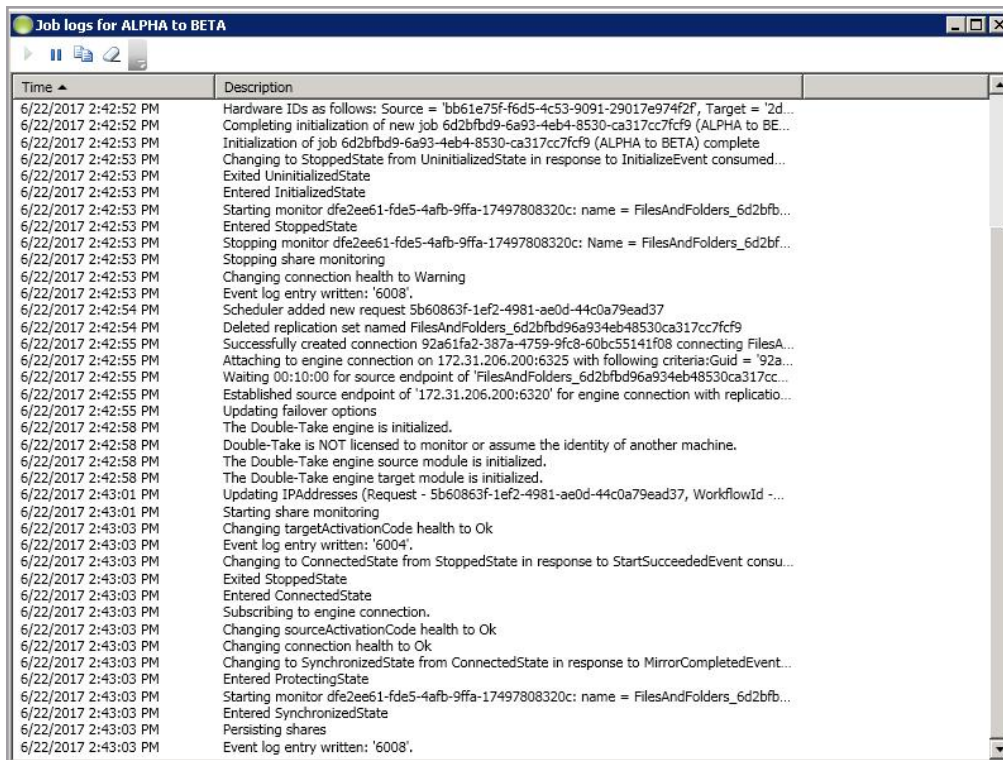
9. Once your servers have passed validation and you are ready to update your job, click **Finish**.

Viewing a full server to ESX job log

You can view a job log file through the 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.



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



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_6d2bf...
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.

Failing over full server to ESX jobs

When a failover condition has been met, failover will be triggered automatically if you disabled the wait for user option during your failover configuration. If the wait for user before failover option is enabled, you will be notified in the console when a failover condition has been met. At that time, you will need to trigger it manually from the console when you are ready.



If you have paused your target, failover will not start if configured for automatic failover, and it cannot be initiated if configured for manual intervention. You must resume the target before failover will automatically start or before you can manually start it.

1. On the **Jobs** page, highlight the job that you want to failover and click **Failover or Cutover** in the toolbar.
2. Select the type of failover to perform.
 - **Failover to live data**—Select this option to initiate a full, live failover using the current data on the target. This option will shutdown the source machine (if it is online), stop the protection job, and start the replica virtual machine on the target with full network connectivity.
 - **Perform test failover**—Select this option to perform a test failover.
 - 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.
 - Scheduled snapshots will be deferred until the test replica is online.
 - 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, including your test failover snapshot. (Performing a live failover or failing over to a snapshot will automatically undo any test in progress.)
 - You can delete the test failover snapshot, if desired, using the **Manage Snapshots** option on the **Jobs** page.
 - **Failover to a snapshot**—Select this option to initiate a full, live failover without using the current data on the target. Instead, select a snapshot and the data on the target will be reverted to that snapshot. This option will not be available if there are no snapshots on the target. This option is also not applicable to clustered environments. To help you understand

what snapshots are available, the **Type** indicates the kind of snapshot.

- **Scheduled**—This snapshot was taken as part of a periodic snapshot.
 - **Deferred**—This snapshot was taken as part of a periodic snapshot, although it did not occur at the specified interval because the job between the source and target was not in a good state.
 - **Manual**—This snapshot was taken manually by a user.
3. Select how you want to handle the data in the target queue.
- **Apply data in target queues before failover or cutover**—All of the data in the target queue will be applied before failover begins. The advantage to this option is that all of the data that the target has received will be applied before failover begins. The disadvantage to this option is depending on the amount of data in queue, the amount of time to apply all of the data could be lengthy.
 - **Discard data in the target queues and failover or cutover immediately**—All of the data in the target queue will be discarded and failover will begin immediately. The advantage to this option is that failover will occur immediately. The disadvantage is that any data in the target queue will be lost.
 - **Revert to last good snapshot if target data state is bad**—If the target data is in a bad state, Carbonite Availability will automatically revert to the last good Carbonite Availability snapshot before failover begins. If the target data is in a good state, Carbonite Availability will not revert the target data. Instead, Carbonite Availability will apply the data in the target queue and then failover. The advantage to this option is that good data on the target is guaranteed to be used. The disadvantage is that if the target data state is bad, you will lose any data between the last good snapshot and the failure.
4. When you are ready to begin failover, click **Failover**.



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

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

- **Retail licensing**—Retail licensing allows the activation of a single operating system installation.
 1. Open the **System** applet in Windows **Control Panel**.
 2. Under **Windows activation** at the bottom of the page, click **Change product key**.
 3. Enter your retail license key. You may need access to the Internet or to call Microsoft to complete the activation.
- **MAK volume licensing**—Multiple Activation Key (MAK) licensing allows the activation of multiple operating system installations using the same activation key.

1. View or download the Microsoft Volume Activation Deployment Guide from the Microsoft web site.
 2. Using an administrative user account, open a command prompt and follow the instructions from the deployment guide to activate MAK clients. Multiple reboots may be necessary before you can access a command prompt. You may need access to the Internet or to call Microsoft to complete the activation.
- **KMS volume licensing**—Key Management Service (KMS) licensing allows IT professionals to complete activations on their local network without contacting Microsoft.
 1. View or download the Microsoft Volume Activation Deployment Guide from the Microsoft web site.
 2. Using an administrative user account, open a command prompt and follow the instructions from the deployment guide to convert a MAK activation client to a KMS client. Multiple reboots may be necessary before you can access a command prompt.

Depending on your replica configuration, you may have to reboot your replica after failover. You will be prompted to reboot if it is necessary.

After failover, if you attempt to use a full server job to revert back to your original configuration, you will need to perform a few additional tasks before creating the full server job. Contact technical support if you need assistance with these steps.

1. On either the source or target, stop the Double-Take and Double-Take Management Service services.
2. Remove the GUID value from HKEY_LOCAL_MACHINE\SOFTWARE\NSI Software\Double-Take\CurrentVersion\Communication\Uniqueld. Do not delete the Uniqueld key. Only delete the GUI value within the key.
3. Restart the the Double-Take and Double-Take Management Service services.
4. Remove and then add your servers back into the Carbonite Replication Console.
5. Install a different license on the original source and complete a host transfer if necessary.

If your job was using vCenter, you may have problems with failover if vCenter is down or if it is unreachable. Contact technical support for details on how to complete failover 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, one or more 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 or remove these hidden NICs.

5. If you performed a test failover, you can undo it by selecting **Undo Failover or Cutover** in the toolbar. Confirm the undo process when prompted. The replica virtual machine on the target will be shut down and, if configured, the virtual disks used for the test failover will be deleted.
6. If you performed a live or snapshot failover, you can reverse your job to a virtual machine running on an ESX host. A reverse job takes the replica virtual machine you failed over to and creates a new protection job from that replica virtual machine to a new replica virtual machine on an ESX host. You cannot reverse to a physical server or non-ESX hosted virtual machine. If you want to reverse to a physical server or non-ESX hosted virtual machine, you will need to delete the job and manually create a new job. See *Reversing protection after failover for full server to ESX jobs* on page 364 for details on reversing a job and creating a new job.

Reversing protection after failover for full server to ESX jobs

If you performed a live or snapshot failover, you can reverse your job to a virtual machine running on an ESX host. A reverse job takes the replica virtual machine you failed over to and creates a new protection job from that replica virtual machine to a new replica virtual machine on an ESX host.

1. If necessary, edit your job to configure reverse job settings. See *Editing a full server to ESX job* on page 356 for details on editing a job.
2. On the **Jobs** page, highlight the failed over job and click **Reverse**. If you just failed over, it may take several minutes after the replica server is online before the **Reverse** option is available.
3. Confirm you want to create the reverse job.

After the reverse job is created, you will be able to control it like any other job. See *Managing and controlling full server to ESX jobs* on page 341.

You cannot reverse to a physical server or non-ESX hosted virtual machine. If you want to reverse to a physical server or non-ESX hosted virtual machine, you will need to delete the job and manually create a new job. The process will depend on where you want to create the new job.

- **Physical server**—Use these steps if you want to create a job to a physical server.
 1. If you are using your original source, resolve the problems on the original source that caused it to fail. If you need to deploy a new server, use the same operating system and disk configuration as the original source.
 2. If Carbonite Availability is still running on the source, replace the license since that license is currently running on the failed over server. If Carbonite Availability is not installed, install it with an appropriate license.
 3. Create a full server job from the failed over replica server to your physical server. See *Creating a full server job* on page 149 for details on creating this job.
 4. Once the initial mirror is complete, failover the full server job. See *Failing over full server jobs* on page 192 for details on this process.
- **Hyper-V virtual server**—Use these steps if want to create a job to a virtual server on a Hyper-V host.
 1. If you are using your original source, delete it from the Hyper-V host. If you want to reuse the .vhd files again, only delete the virtual server from the Hyper-V inventory.
 2. If it is not already, install and license Carbonite Availability on the Hyper-V host you will be using. The host will be the target of the new job you are going to create.
 3. Create a full server to Hyper-V job from your failed over replica server to the Hyper-V host. See *Creating a full server to Hyper-V job* on page 261 for details on creating this job.
 4. Once the initial mirror is complete, failover the full server to Hyper-V job. See *Failing over full server to Hyper-V jobs* on page 302 for details on this process.

Chapter 10 Simulating protection

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

Before and after simulating a job, you should gather network and system information specific to Carbonite Availability operations. Use the Carbonite Replication Console to automatically collect this data. It gathers Carbonite Availability log files; Carbonite Availability and system settings; network configuration information such as IP, WINS and DNS addresses; and other data which may be necessary in evaluating Carbonite Availability 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 protection job, noting the following caveats.
 - When selecting your target, select the **Diagnostics job** checkbox instead of a target server.
 - When you get to the **Set Options** page in the workflow, some options for your selected job type will not be displayed because they are not applicable. For example, target specific selections will not be displayed because there is no actual target with the TDU.
3. Once you have established your job, you should ideally let it run for several days to gather accurate data for your network and source server usage. The simulation data will be logged to the Carbonite Availability 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 11 Special network configurations

Carbonite Availability 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 367
- *IP and port forwarding* on page 368
- *Domain controllers* on page 371
- *NetBIOS* on page 372
- *WINS* on page 373
- *DNS* on page 375
- *Non-Microsoft DNS* on page 383
- *Macintosh shares* on page 385
- *NFS Shares* on page 386

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 Availability ports**—Ports 6320, 6325, and 6326 are used for Carbonite Availability communications and must be open on your firewall. Open TCP for both inbound and outbound traffic. Carbonite Availability uses ICMP pings, by default, to monitor the source for failover. You should configure your hardware to allow ICMP pings between the source and target. If you cannot, you will have to configure Carbonite Availability to monitor for a failure using the Double-Take service. See the failover instructions for your specific job type.
- **Microsoft WMI and RPC ports**—Some features of Carbonite Availability 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 Availability 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 Availability 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 Availability 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: [masked]

Domain:

Management Service port: 1025 Use default port

Add

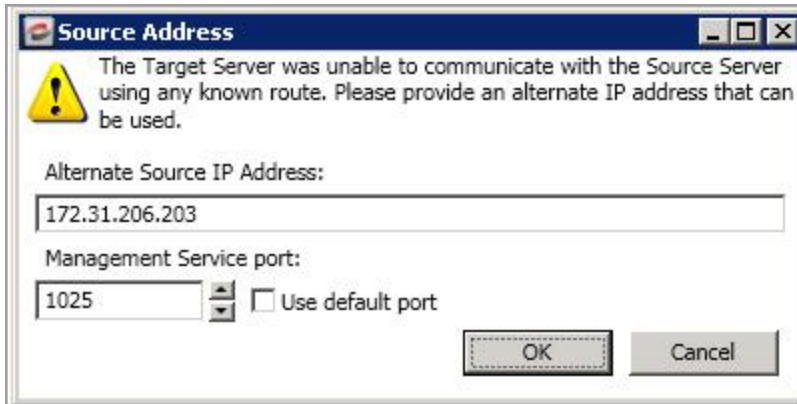
Servers to be added:

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 failover. 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
 Change target ports to match source during failover

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 this target IP address: Management Service port: Use default port Replication Service port: Use default port

Receive commands on this source IP address: Use default route

Reverse Protection and Routing

Send data to this target IP address:

Management Service port: Use default port

Replication Service port: Use default port

Receive commands on this source IP address: 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:

Select a reserved IP address on the target:

Domain controllers

Failover of domain controllers is dependent on the Carbonite Availability functionality you are using.

- **Domain controller role**—If you want to failover a domain controller, including the roles of the domain controller, you should create full server or virtual protection.
- **Non-domain controller role**—If you are only protecting data, you can failover a domain controller, but the roles of the domain controller are not included. The server will be a member server after failover. In this case, you need to keep in mind that the unavailability of some of the FSMO (Flexible Single Master Operation) roles can cause immediate issues such as the inability to extend the Active Directory schema or to add a domain to a forest.
- **Global catalog server**—If your source is a global catalog server, you should have other global catalog servers throughout the network to ensure that the failure of the source will not impact users.

NetBIOS

Because NetBIOS is not available on Windows 2008 R2 or later, if you are using a workgroup environment and do not have DNS host records, there may be a delay of up to five minutes before the failed over source server name is available for client access. See *About the NetBIOS Interface* in the [MSDN library](#).

WINS

When Carbonite Availability failover occurs, Windows initiates WINS registration with the target's primary WINS server to associate the source server's name with the target's primary IP address. In an environment with just one WINS server, no additional processing is required. In an environment with more than one WINS server, WINS replication will distribute the updated WINS registration to other WINS servers on the network. The length of time required for all WINS servers to obtain the new registration depends on the number of WINS servers, the WINS replication architecture, and the WINS replication interval. Clients will be unable to access the target until their WINS server has received the updated WINS information. You can reduce the time required for the WINS updates, thereby decreasing the wait time for the end users, by scripting the WINS updates in the Carbonite Availability failover scripts. You have two options for scripting the WINS updates.

- *WINS registration* on page 373—This option registers a user-specified server with WINS. It requires less network overhead but administrator group membership on all WINS servers.
- *WINS replication* on page 374—This option forces WINS replication. It does not require any special privileges, but requires system and network resources to complete WINS replication. The impact on the network will depend on the size and complexity of the WINS architecture.

WINS registration

WINS registration can be added to your failover and failback scripts by using the Windows NETSH command with the WINS add name context. Add the following command to your failover and failback scripts for each additional WINS server in your environment (excluding the target's primary WINS server).

```
netsh wins server wins_server_IP_address add name Name=source_server_name RecType=1  
IP={IP_address}
```

Use the following variable substitutions.

- *wins_server_IP_address*—The IP address of the WINS server
- *source_server_name*—The name of the source server
- *IP_address*—The IP address of the target that has taken over for the failed source (for the failover script) or the IP address of the source that is reassuming its original identity (for the failback script)

For example, suppose you had the following environment.

- **Source name and IP address**—Alpha 192.168.1.108
- **Target name and IP address**—Beta 116.123.2.47
- **Target's Primary WINS server**—116.123.2.50
- **First secondary WINS server on the network**—192.168.1.110
- **Second secondary WINS server on the network**—150.172.114.74

You would add the following to your failover script to register the source's name with the target's IP address on the two secondary WINS servers.

```
netsh wins server 192.168.1.110 add name Name=Alpha RecType=1 IP={116.123.2.47}
netsh wins server 150.172.114.74 add name Name=Alpha RecType=1 IP={116.123.2.47}
```

You would add the following to your failback script to register the source's name back with the source's original IP address on the two secondary WINS servers.

```
netsh wins server 192.168.1.110 add name Name=Alpha RecType=1 IP={192.168.1.108}
netsh wins server 150.172.114.74 add name Name=Alpha RecType=1 IP={192.168.1.108}
```

See your Windows documentation or the Microsoft web site for more details on the NETSH command.

WINS replication

WINS replication can be added to your failover and failback scripts by using the Windows NETSH command with the WINS set replicate context. Add the following command to your failover and failback scripts.

```
netsh wins server target's_primary_wins_server_IP_address set replicateflag 1
```

Use the following variable substitution.

- ***target's_primary_wins_server_IP_address***—The IP address of the target's primary WINS server

For example, suppose you had the following environment.

- **Source name and IP address**—Alpha 192.168.1.108
- **Target name and IP address**—Beta 116.123.2.47
- **Target's Primary WINS server**—116.123.2.50
- **First secondary WINS server on the network**—192.168.1.110
- **Second secondary WINS server on the network**—150.172.114.74

You would add the following to your failover script to force the target's primary WINS server to replicate its updated information to the other secondary WINS servers on the network.

```
netsh wins server 116.123.2.50 set replicateflag 1
```

You would add the same line to your failback script to force the target's primary WINS server to replicate its updated information again. This would replicate information for the source's name and the source's original IP address to the other secondary WINS servers on the network.

```
netsh wins server 116.123.2.50 set replicateflag 1
```

See your Windows documentation or the Microsoft web site for more details on the NETSH command.

DNS

If you are using a Microsoft DNS server, when Carbonite Availability failover occurs, DNS may or may not be automatically updated depending on your job type and job options. If the end-users use DNS to resolve server names and the source IP address was not failed over to the target, additional DNS updates will be required because the host records for the source will remain intact after failover. You can automate this process by scripting the DNS updates in the failover and failback scripts. You have two options for scripting the DNS updates.

- *Windows DNSCMD command* on page 375—The Windows Support Tools contain a DNS Server Troubleshooting Tool utility. This utility includes the DNSCMD command which can be scripted to delete and add host and reverse lookup entries in DNS.
- *Carbonite Availability DFO utility* on page 377—Carbonite Availability also has a utility, called DFO (DNS Failover). The DFO utility can be used to script the deletion and addition of the host and reverse lookup entries in DNS. This utility is installed with Carbonite Availability.

Windows DNSCMD command

DNS updates can be added to your failover and failback scripts by using the Windows DNSCMD command as long as dynamic updates are enabled on the DNS zone and the account running the Double-Take service is a member of the DNSAdmins security group. (See your Microsoft documentation to verify if dynamic updates are enabled.) You may want to disable the DNS registration feature of each IP address that is being changed in DNS to prevent the source from changing the record back when it comes online after a failover.

Add the following commands to your failover and failback scripts to delete the host and reverse lookup entries and add new entries associating the source to the target.

- `dnscmd DNS_server's_FQDN /RecordDelete DNS_zone source_server_name A source_server_IP_address /f`
- `dnscmd DNS_server's_FQDN /RecordDelete www.xxx.in-addr.arpa zzz.yyy PTR source_server's_FQDN /f`
- `dnscmd DNS_server's_FQDN /RecordAdd DNS_zone source_server_name A target_server_IP_address`
- `dnscmd DNS_server's_FQDN /RecordAdd aaa.bbb.in-addr.arpa ddd.ccc PTR source_server's_FQDN`

Use the following variable substitutions.

- `DNS_server's_FQDN`—The fully qualified domain name of the DNS server
- `DNS_zone`—The name of the DNS zone
- `source_server_name`—The name of the source server
- `source_server_IP_address`—The IP address on the source
- `www.xxx`—The first two octets of the source's IP address. For example, if the source's IP address is 192.168.1.108, this variable would be 192.168.
- `zzz.yyy`—The last two octets, in reverse order, of the source's IP address. For example, if the source's IP address is 192.168.1.108, this variable would be 108.1.
- `source_server's_FQDN`—The fully qualified domain name of the source server

- `target_server_IP_address`—The IP address on the target
- `aaa.bbb`—The first two octets of the target's IP address. For example, if the target's IP address is 116.123.2.47, this variable would be 116.123.
- `ddd.ccc`—The last two octets, in reverse order, of the target's IP address. For example, if the target's IP address is 116.123.2.47, this variable would be 47.2.

For example, suppose you had the following environment.

- Full qualified domain name of the source—Alpha.domain.com
- Source IP address—192.168.1.108
- Fully qualified domain name of the target—Beta.domain.com
- Target IP address—116.123.2.47
- Fully qualified domain name of the DNS server—DNSServer.domain.com
- DNS zone—domain.com

You would add the following to your failover script to delete the host and reverse lookup entries and add new entries associating the source to the target.

```
dnscmd DNSServer.domain.com /RecordDelete domain.com alpha A 192.168.1.108 /f
dnscmd DNSServer.domain.com /RecordDelete 192.168.in-addr.arpa 108.1 PTR alpha.domain.com /f
dnscmd DNSServer.domain.com /RecordAdd domain.com alpha A 116.123.2.47
dnscmd DNSServer.domain.com /RecordAdd 116.123.in-addr.arpa 47.2 PTR alpha.domain.com
```

You would add the following to your failback script to delete the host and reverse lookup entries and add new entries associating the source with its original identity.

```
dnscmd DNSServer.domain.com /RecordDelete domain.com alpha A 116.123.2.47 /f
dnscmd DNSServer.domain.com /RecordDelete 116.123.in-addr.arpa 47.2 PTR alpha.domain.com /f
dnscmd DNSServer.domain.com /RecordAdd domain.com alpha A 192.168.1.108
dnscmd DNSServer.domain.com /RecordAdd 192.168.in-addr.arpa 108.1 PTR alpha.domain.com
```

See your Windows documentation or the Microsoft web site for more details on the DNSCMD command.

Carbonite Availability DFO utility

DNS updates can be added to your failover and failback scripts by using the Carbonite Availability DFO utility as long as the utility has been registered and the proper privileges are configured.

How the DFO utility works

The DFO utility performs DNS resource record modifications by connecting to the DNS namespace (root\microsoftdns) on the DNS server using WMI. The WMI connection can be made using passed credentials or impersonation if the account running the DFO utility has permissions to perform all DNS-related activities. Passed credentials can be encrypted using Microsoft's CAPICOM dynamic link library with DFO specifying the triple DES encryption algorithm with the maximum key length available (168). By providing reliable encryption, the DFO utility allows you to avoid storing secure passwords in script files.

If the source experiences a failure or an extended outage, clients will need to be redirected automatically to the target server. In these cases, the DFO utility can help make the network redirection portion of failover transparent to end users.

The DFO utility is able to modify five DNS resource record types: A, AAAA, CNAME, MX, and PTR. Here is how it works for the host record or A type.

1. The DFO utility builds and executes a focused WMI query to retrieve a collection of matching source DNS resource records from the DNS server. For example: `SELECT * FROM MicrosoftDNS_AType WHERE IPAddress="192.168.1.108"`
2. The DFO utility iterates through the returned collection and modifies any matching resource records.
 - a. The DFO utility spawns an instance of the WMI DNS resource record object to call the modify method.
 - b. The DFO utility sets the parameters such that the target IP address is the input parameter.
 - c. The DFO utility executes the modify method on the WMI DNS record object.
3. The DFO utility locks the DNS resource record in Active Directory so that the source computer account is unable to modify the record outside of the DFO utility. This is done to prevent modification of the resource record to point back to the source machine until the failback process is initiated by the user through the DFO utility.
 - a. The DFO utility denies permission to modify the Active Directory object representing the DNS entry.
 - b. The DFO utility gets the DNS resource record object in Active Directory.
 - c. The DFO utility reads the security descriptor and gets the DACL.
 - d. The DFO utility adds the ACE "ACCESS_DENIED" type for the passed-in trustee name (for example, source computer account, cluster administrator account, and so on) to deny access to the "Write All Properties" permission.
4. The DFO utility logs the results of the actions performed/attempted.

Other record types require different queries and input parameters. Additionally, CNAME, MX, and PTR record types do not execute the Active Directory object locking routines that A and AAAA type records require for failover.

- **AAAA type**—Except for the query difference, this record type is identical to the A type record.

```
SELECT * FROM MicrosoftDNS_AAAAType WHERE  
IPAddress="21DA:D3:0:2F3B:2AA:FF:FE28:9C5A"
```
- **CNAME type**—This type does not have Active Directory object locking to prevent updates during failover.

```
SELECT * from MicrosoftDNS_CNAMEType WHERE PrimaryName="sql1.doubletake.com"
```
- **MX type**—This type does not have Active Directory object locking to prevent updates during failover.

```
SELECT * from MicrosoftDNS_MXType WHERE MailExchange="mail1.doubletake.com"
```
- **PTR type**—Instead of modifying the source record, the PTR type deletes the source PTR record and create a new PTR record by using previous source PTR text record information, substituting the target FQDN for the source FQDN, and calling the CreateInstanceFromPropertyData() method on the DNS server. This type does not have Active Directory object locking to prevent updates during failover.

```
SELECT * from MicrosoftDNS_PTRType WHERE PTRDomainName="sql1.doubletake.com"
```

During failback, the same mechanisms that were used during failover are used, except that the original source-related records are modified to point to the original source. (During failover, the source records were modified to point to the target IP address or name, depending on the record type.) Also, during failover the A and AAAA type DNS resource records are modified in DNS and then locked in Active Directory; during failback, those record types are unlocked in Active Directory and then modified in DNS.

Using the DFO utility

1. From a command prompt, change to the Carbonite Availability program files directory and register the DFO utility by entering the command `regsvr32 capicom.dll`
2. Create a user account that has full control on the WMI DNS namespace on the source's primary DNS server.
 - a. From a command prompt, enter the command `wmimgmt.msc`.
 - b. Right-click **WMI Control** and select **Properties**.
 - c. On the **Security** tab, expand the tree under **Root**.
 - d. Select **MicrosoftDNS** and click **Security**.
 - e. Click **Add** and identify the user account that you want the DFO utility to use.
 - f. Grant the user account permissions for Execute Methods, Enable Account, Remote Enable, and Read Security.
 - g. Click **Advanced** and in the **Permissions** list, select the user account and click **Edit**. Select **This namespace and subnamespaces**.
 - h. Click **OK** to close all open dialog boxes and then close the console.
 - i. Restart the Windows Management Instrumentation service for the changes to take effect.
 - j. From a command prompt, enter the command `dcomcnfg`.
 - k. Expand **Component Services**, expand **Computers**, then right-click **My Computer** and select **Properties**.
 - l. On the **COM Security** tab, under **Access Permissions**, click **Edit Limits**.

- m. Click **Add**, identify the user account, and click **OK**.
 - n. In the **Permissions for User** list, allow permissions for Local Access and Remote Access and click **OK**.
 - o. Under **Launch and Activation Permissions**, click **Edit Limits**.
 - p. Click **Add**, identify the user account, and click **OK**.
 - q. In the **Permissions for User** list, allow permissions for Local Launch, Remote Launch, Local Activation, and Remote Activation.
 - r. Click **OK**.
 - s. Click **OK** again.
 - t. Expand **My Computer**, expand **DCOM Config**, then right-click **Windows Management and Instrumentation** and select **Properties**.
 - u. On the **Security** tab, under **Access Permissions**, click **Edit**.
 - v. Click **Add**, identify the user account, and click **OK**.
 - w. In the **Permissions for User** list, allow permissions for Local Access and Remote Access and click **OK**.
 - x. Click **OK** to close all open dialog boxes.
 - y. Restart the DNS/Domain Controller.
3. Add the same user account that has full control on the WMI DNS namespace to the domain's DnsAdmins group where the source's primary DNS server is located.
 - a. Select **Active Directory Users and Computers** from Administrative Tools.
 - b. Right-click the **DnsAdmins** group and select **Properties**.
 - c. Select the **Members** tab, click **Add**, and identify the user account that you granted full control on the WMI DNS namespace.
 - d. Click **OK** to close all open dialog boxes and then close Active Directory Users and Computers.
 4. Add a user to the Server Operator group.
 - a. Select **Active Directory Users and Computers** from Administrative Tools.
 - b. Select **Builtin**, then right-click the **Server Operators** group and select **Properties**.
 - c. Select the **Members** tab, click **Add**, and identify the user account that you granted full control on the WMI DNS namespace.
 - d. Click **OK** to close all open dialog boxes and then close Active Directory Users and Computers.
 5. Grant the user full control over the source and target DNS records.
 - a. Select **DNS** from Administrative Tools.
 - b. Locate both the source and target records in the forward and reverse lookup zones.
 - c. For each record, right-click and select **Properties**.
 - d. On the **Security** tab, click **Add** and identify the user account that you granted full control on the WMI DNS namespace, and click **OK**.
 - e. In the **Permissions for User** list, allow permissions for **Full control** and click **OK**.
 - f. Click **OK** to close all open dialog boxes and repeat for each record.
 - g. Close DNS Manager.
 6. Add the appropriate DFO command to your failover script using the following syntax.

Command

DFO

Description

Used in scripts to failover DNS server name

Syntax

```
DFO [/DNSSRVNAME <dns_server_name>] [/SRCNAME <source_fqd_name>] [/SRCIP <source_ip>] [/TARIP <target_ip>] [/TARNAME <target_fqd_name>] [/RECORDTYPE <rec_type>] [/USERNAME <user_name>] [/PASSWORD <password>] [/DNSZONE <zone_name>] [/DNSDOMAIN <domain_name>] [/LOGFILE <file_name>] [/FAILBACK [fb_switch]] [/SETPASSWORD <user_name> <password>[machine][file]] [/GETPASSWORD] [/LOCK] [/UNLOCK] [/TRUSTEE [<trustee_name>]] [/VERBOSE] [/FLUSHDNS] [/MACHINE <machine_fqd_name>] [/TTL <seconds>] [/ADDDOMAIN <active_directory_domain_name>] [/SOURCEDN <source_domain_name>] [/TEST] [/DEBUG] [/HELP]
```

Options

- **DNSSRVNAME** dns_server_name—The name of the source domain/zone's primary DNS server. If not specified, the local machine will be used.
- **SRCNAME** source_fqd_name—The source machine's fully qualified domain name
- **SRCIP** source_ip—The source machine's IP address
- **TARIP** target_ip—The target machine's IP address
- **TARNAME** target_fqd_name—The target machine's fully qualified domain name (required only for failback)
- **RECORDTYPE** rec_type—The type of DNS resource records to modify or list. Values record types are ALL, MSEXCHANGE, A, CNAME, MX, PTR, STD, or STANDARD. STD and STANDARD are used to specify a non-Exchange record (minus the MX records). By default, all record types are included.
- **USERNAME** user_name—The domain name of the user account. If not specified, the account running the utility will be used.
- **PASSWORD** password—The password associated with the user account
- **DNSZONE** zone_name—The name of the DNS zone or DNS container, used to refine queries
- **DNSDOMAIN** domain_name—The name of the DNS domain, used to refine queries
- **LOGFILE** file_name—The name of the log file
- **FAILBACK** fb_switch—Denotes a failback procedure, performed after a failed source is recovered or restored (required for failback). By default, the DFO will only failback records in the dfo_failback_config.dat file. The fb_switch is optional and allows you to enter search criteria to identify the records

to change back, even if they are not in the configuration file. The `fb_switch` is also used if the `dfo_failback_config.dat` file is missing.

- **SETPASSWORD** `user_name password machine file`—Stores user credentials on the specified machine or in the specified file for later use. The file will be encrypted. This option must be run separately from a modify or list activity.
- **GETPASSWORD**—Retrieves previously stored user credentials. This option can only be used if the credentials were previously stored with the `setpassword` option.
- **LOCK**—Allows Active Directory locking for the A record type of the source specified without modifying the record
- **UNLOCK**—Allows Active Directory unlocking for the A record type of the source specified without modifying the record
- **TRUSTEE** `trustee_name`—The domain account for the source machine (`domain\machine$`). DFO attempts to deny write permissions to the DNS A record on failover for the account identified as the trustee. “Deny write permissions” is then removed from the DNS A record on failback. This keeps the source server from reclaiming its DNS A record if it comes back online prior to failback.
- **VERBOSE**—Logging and display level set to maximum detail
- **FLUSHDNS**—Runs the `ipconfig /flushdns` command to flush the DNS cache.
- **MACHINE** `machine_fqd_name`—Specifies the machine where `ipconfig /flushdns` is run. Use the fully-qualified domain name of the machine.
- **TTL** `seconds`—Specifies the number of seconds for the time to live value of all modified records
- **ADDDOMAIN** `active_directory_domain_name`—The name of the Active Directory domain
- **SOURCEDN** `source_domain_name`—The name of the source's domain
- **TEST**—Runs in test mode so that modifications are not made, only listed
- **DEBUG**—Forces DFO to write the DNS resource record as-is to the `dfolog.log` file prior to any DFO modify or list activity.
- **HELP**—Displays the syntax of the DNS Failover utility

Examples

- `dfo /dnssrvname gamma.domain.com /srcname alpha.domain.com /srcip 206.31.4.10 /verbose` (Lists all resource records on the specified DNS server that match the source criteria)
- `dfo /dnssrvname gamma.domain.com /srcname alpha.domain.com /srcip 206.31.4.10 /tarip 210.11.12.13 /verbose` (Modifies all resource records on the specified DNS server that match the source criteria, using the credentials of the account running the utility to connect to the DNS server)
- `dfo /dnssrvname gamma.domain.com /srcname alpha.domain.com /srcip 206.31.4.10 /tarip 210.11.12.13 /username domain.com\admin /password /verbose` (Modifies all resource records on the specified DNS server that match the source criteria, using the username and password to connect to the DNS server)

- `dfo /dnssrvname gamma.domain.com /srcname alpha.domain.com /scip 210.11.12.13 /tarname beta.domain.com /tarip 206.31.4.10 /failback /verbose` (Fails back all resource records on the specified DNS server that were changed on failover)
- `dfo /setpassword domain.com\admin password` (Stores the user name and password in an encrypted file)
- `dfo /dnssrvname gamma.domain.com /srcname alpha.domain.com /scip 206.31.4.10 /tarip 210.11.12.13 /username domain.com\admin /getpassword /verbose` (Modifies all resource records on the specified DNS server that match the source criteria, using the specified username and retrieving the password from the encrypted file)

Notes

All options are marked as optional, enclosed in brackets [], however, you will have to supply options to execute DFO functionality. The options to supply will depend on the functionality you are trying to complete. For example, you must supply the username and password to cache credentials, but you do not need those options to query or modify a DNS record.

Non-Microsoft DNS

If you are using a non-Microsoft DNS server (such as Unix) or if you are in a non-domain configuration (such as a workgroup), when Carbonite Availability failover occurs, DNS may or may not be automatically updated depending on your job type and your job options. If the end-users use DNS to resolve server names and the source IP address was not failed over to the target, additional DNS updates will be required because the host records for the source will remain intact after failover. You can automate this process by scripting the DNS updates in the failover and failback scripts.

One option is to use a BIND DNS client for DNS scripting. The following steps provide an example of how you can use a BIND DNS client for DNS failover and failback scripting. You may need to modify this example to fit your environment.

1. Go to www.isc.org and download the appropriate BIND DNS client.
2. Install the BIND client on the target server.
3. Set a PATH statement for the BIND directory to ensure that it runs every time the executable is called.
4. Create a failover script file in the Carbonite Availability directory.
5. Add the following line to the failover script file, substituting your Carbonite Availability directory for `install_location`.

```
nsupdate.exe "c:\install_location\dnsover.txt"
```

6. Save the failover script file.
7. Create a text file, called `dnsover.txt` in the Carbonite Availability directory.
8. Add the following lines to the `dnsover.txt` file, substituting your source name, fully-qualified domain name, target name, and target IP address as appropriate.

```
update delete source_server_name.fully_qualified_domain_name.com A  
update add target_server_name.fully_qualified_domain_name.com 86400 A target_server_IP_address  
send
```

9. Save the `dnsover.txt` file.
10. Create a failback script file in the Carbonite Availability directory.
11. Add the following line to the failback script file, substituting your Carbonite Availability directory for `install_location`.

```
nsupdate.exe "c:\install_location\dnsback.txt"
```

12. Save the failback script file.
13. Create a text file, called `dnsback.txt` in the Carbonite Availability directory.
14. Add the following lines to the `dnsback.txt` file, substituting your target name, fully-qualified domain name, source name, and source IP address as appropriate.

```
update delete target_server_name.fully_qualified_domain_name.com A  
update add source_server_name.fully_qualified_domain_name.com 86400 A source_server_IP_address  
send
```

15. Save the `dnsback.txt` file.

16. Change the Double-Take service on the target server to a domain account that has rights to modify BIND DNS. Stop and start the service to have it take effect.

Macintosh shares

A share is any volume, drive, or directory resource that is shared across a network. During failover, the target can assume or add any source shares so that they remain accessible to the end users. Automatic share failover only occurs for standard Windows file system shares. Other shares, including Macintosh volumes, must be configured for failover through the failover scripts or created manually on the target.

1. On your target, set the File Server for Macintosh service to manual startup. This allows the post-failover script on the target to control when the service starts on the target.
2. Create each volume on the target machine exactly as it exists on the source. Use the Shared Folder wizard to configure each volume as a Macintosh-accessible volume. Follow these steps to start the wizard.
 - a. Open the Control Panel and click **Administrative Tools**.
 - b. Select **Configure Your Server**.
 - c. In the Configure Your Server window, click the **File Server** link.
 - d. Click **Start the Shared Folder wizard** to start the wizard, and then follow the directions provided by the wizard. On the Create Shared Folders screen, you must enable **Apple Macintosh**.



You can automate the creation of the volumes during the failover process by using the macfile volume command in the post-failover batch file. For detailed information on how to use this command, see your Windows reference guide.

3. On the target machine, copy the chngname utility, chngname.exe, from the \tools\Win2K directory of the Carbonite Availability DVD or from the Carbonite support web site to the directory where Carbonite Availability 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 Availability 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 Availability 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 Availability 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 Availability DVD or from the support web site to the directory where Carbonite Availability 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 Availability 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 Availability chngname utility and replace
rem source_name with the name of the source machine.
<drive>\<directory>\chngname /s source_name
rem The following command starts the NFS service
net start "Server for NFS"
```

In the event of a failure, the clients must remount the shares in order to access them.

Chapter 12 Recommended optimizations

Carbonite Availability is an exceptionally flexible product that can be used in a wide variety of network configurations. However, this flexibility can make implementing Carbonite Availability effectively difficult. There is often a balance that must be found between various configuration options and their relative benefits.

Through years of testing and implementing in diverse environments, Carbonite has compiled the following list of recommended optimizations. Keep in mind, what works for one environment or configuration may not work in another. A best practice in one organization may be ineffective in another. You should work with Carbonite technical support or Professional Services when making optimization changes.

- *Planning* on page 388
- *Installation optimizations* on page 389
- *General optimizations* on page 390
- *Full server optimizations* on page 394
- *Application optimizations* on page 395

Planning

Before you begin your Carbonite Availability installation, you should plan your implementation strategy. Ask yourself the following questions.

- What is the role of each server? Will this server be a source? Will this server be a target?
- Is the source server a Domain Controller? Or does it have another very specific role or configuration? You may want consider protecting the entire server in these cases.
- Is the source running Microsoft SQL?
- How much data will you be protecting? Can your target handle that amount of data?
- How much bandwidth is available between your source and target? Can your network handle the mirroring and replication traffic between the two servers? If the amount of change is greater than the bandwidth available, you may want to consider getting additional bandwidth or planning for disk queuing.

If there are concerns about resource utilization or how Carbonite Availability replication will impact the environment, you can profile the source server, the network links between the source and target, and the target server before installing Carbonite Availability to ensure that each component has adequate resources to handle the added load of replicating the data. Most environments do not require this type of analysis, but it may be needed if there are applications producing high-volume file writes or limited CPU, memory, disk, or network resources.

The best way to understand the impact of replication in an environment is to set up test equipment that simulates the production environment. However, if the resources to test in this manner are not available, resource utilization can be analyzed using Windows Performance Monitor and a utility to monitor network utilization. Performance data should be logged for a period that encompasses normal usage as well as any maintenance, backup, scheduled jobs, or batch processing that occurs. If utilization of any component is extremely high for a significant period of time, then it may be necessary to modify particular Carbonite Availability options. Keep in mind that some factors that are typically not in a test environment, such as backups and other applications using bandwidth, can affect resource utilization in the production environment.

One method to avoid for planning purposes is estimating the amount of data that will be replicated in a given period using the amount of data backed up in a differential backup. Although this may be valid in some cases, it is usually not a good indicator because it is based on the differences in data at the time of backup. For example, if a 1 MB Microsoft Word document is saved ten times throughout the day, this will result in 10 MB of replication traffic because Word rewrites the entire file each time it is saved. However, this will only result in 1 MB being backed up for a differential backup.

Installation optimizations

Make sure you review the requirements for your job type. When you perform the installation, you will have several decisions to make.

- **Login**—Always log on to the server with an account that is in the local Administrators group before starting the installation.
- **Components**—Decide what components to install and where to install them. Keep in mind that server components are required for systems that will function as a source or target, and they require a license key for the service to run. Client components do not require a license key, but are required to administer Carbonite Availability servers throughout your environment.
- **License key**—The license key that is required for each server is a 24-character, alpha-numeric key which applies the appropriate license to your installation.
- **Queues**—The installation will prompt you to select disk queue settings. Carbonite Availability uses system memory to store data. When the Carbonite Availability system memory limit is reached, Carbonite Availability will queue to disk.
 - If you set the system memory limit lower, Carbonite Availability will use less system memory, but you will queue to disk sooner which may impact system performance. If you set it higher, Carbonite Availability will maximize system performance by not queuing to disk as soon, but the system may have to swap the memory to disk if the system memory is not available. In general, the amount of memory Carbonite Availability and other applications on the server are configured to use should be less than the amount of physical memory on the system to prevent low memory conditions.
 - Select your disk queue location for optimal performance. For example, do not put it on the same physical device as the data being replicated. If possible, put it on a dedicated array optimized for writing. If you expect large amounts of disk queuing, you may want to increase the size of the queue files from the default of 5 MB to 50 MB for more efficient queuing.

See *Carbonite Availability queue* on page 57 for more details on the disk queue usage.

- **Upgrades**—Keep the following caveats in mind when upgrading.
 - If Carbonite Availability does not function correctly after the upgrade, run the Carbonite Availability Setup, select the **Repair** option, and reboot the server. If Carbonite Availability does not function correctly after the repair, uninstall Carbonite Availability, reboot, and install the new version.
 - If your current Carbonite Availability version is more than two minor versions old, you may want to consider uninstalling the old version and installing the new version instead of upgrading.
 - Always upgrade the target server first when upgrading a source and target configuration.

General optimizations

The following are general optimizations that can be used for any Carbonite Availability job type.

- *Performance optimizations* on page 390
- *General manageability* on page 392
- *Anti-virus protection* on page 393
- *Hardware configurations* on page 393

Performance optimizations

- **Initial mirror across slow network**—A large amount of data that is being mirrored across a slow network may take days to complete the initial mirror, depending on the amount of data and the available bandwidth. You may want to consider the following options to reduce the amount of time for the initial mirror.
 - Move the target server to the source's site for the initial mirror. When the mirror is complete, delete the job, move the target server to its permanent location and create a new job using a difference mirror.
 - Archive the data to media that can be transported to the target site and restored to the target server. When the data is restored to the target, create the job using a difference mirror.
 - Create a compressed archive of the source data, copy the archive to the target, decompress the data, and then create the job using a difference mirror.
- **Compression**—Carbonite Availability compression should be used when network bandwidth between the source and target is limited. In some cases, performance may also be improved by enabling compression in high-bandwidth environments. The best level of compression for a given solution will depend on a number of factors, including the type of data being replicated, CPU load, and available bandwidth. Since compression settings can be changed dynamically, the easiest way to find the best level of compression is to enable the mid-level and monitor the results. If data is still being queued on the source, increase the compression level. If CPU load becomes an issue on the server, decrease the compression level.
- **Low bandwidth and queuing**—In low bandwidth environments, you may need to revisit the queuing configuration you established when you were installing Carbonite Availability. See the *Installation optimizations* on page 389 and *Carbonite Availability queue* on page 57 for more details on the disk queue usage.
- **High latency and mirror packet size**—In a high latency environment (greater than 100 ms response times), you may want to consider increasing the size of the packets of mirror data. The default value is 65536 bytes. You may want to double that to 131072 bytes. However, if the average size of the files on the source is smaller than the value you set, changing the value will not help. This option is available through the *Source server properties* on page 61.
- **High latency and MaxChecksumBlocks**—In a high latency environment (greater than 100 ms response times), you may want to consider increasing the number of checksum values retrieved from the target. The default is 32. You may want to double that to 64. See the *MaxChecksumBlocks* server setting in the *Reference Guide*.
- **Target write speed**—In high-bandwidth environments, Carbonite Availability throughput is most often limited by the write speed of the target disks. Accordingly, optimizing the target disks for write performance will often increase Carbonite Availability performance, particularly for full

mirrors and high loads of replication. Using RAID 0 and/or RAID 1 instead of RAID 5 on the target disks will improve the target write performance, as well as allocating some (or all) of the I/O controller's cache memory to write operations.

- **TCPBufferSize**—Network throughput is directly related to the TCP buffer size and the network latency of the LAN or WAN connection. By default, Carbonite Availability is configured for a 1Gbit LAN network. If you are replicating across a different LAN network or a WAN network, adjust the TCP buffer size accordingly. For example, for a 100Mbit LAN, the value should be around 37500, and for a WAN, the value should be around 130000. See the TCPBufferSize server setting in the *Reference Guide*.
- **Windows MTU**—The Maximum Transmission Unit (MTU) is the largest amount of data, a packet, that can be transferred in one physical frame on a network. If the MTU is too high, you may get fragmented packets which can slow down Carbonite Availability mirroring and replication and can possibly cause lost Carbonite Availability connections. Use the ping command with the -f -l 1500 options. If you receive a response that packets need to be fragmented, you should lower your MTU value. See the Microsoft article [314825](#) for details on specifying the MTU value.
- **Disable root encryption**—If the top-level folders in your jobs are not encrypted, you can gain a performance improvement by disabling root encryption. See the EnableRootEncryption server setting in the *Reference Guide*.

General manageability

- **Temporary files**—Some applications create temporary files that are used to store information that may not be necessary to replicate. If user profiles and home directories are stored on a server and replicated, some unexpected data may be replicated if applications use the \Local Settings\Temp directory to store data. This could result in significant amount of unnecessary data replication on large file servers. Additionally, the \Local Settings\Temporary Internet Files or \AppData\Local\Microsoft\Windows\Temporary Internet Files directories can easily reach a few thousand files and dozens of megabytes. When this is multiplied by a hundred users it can quickly add up to several gigabytes of data that do not need to be replicated. You may want to consider excluding temporary data like this, however it is important to know how applications may use these temporary files. For example, Microsoft Word creates a temporary file when a document is opened. When the user closes the file, the temporary file is renamed to the original file and the original file is deleted. In this case, you must replicate that temporary file so that Carbonite Availability can process the rename and delete operations appropriately on the target.
- **E-mail notification**—Enable e-mail notification through the *E-mail notification configuration* on page 65 so that you are notified when a Carbonite Availability message is written to the Event log for that server.
- **Target path blocking**—Target path blocking prevents the modification of the copy of the source data on the target until failover has occurred or protection is disabled. This can be configured for some job types or for all jobs to a target through the *Target server properties* on page 63.
- **Disable attribute replication**—On servers where the file permissions need to be different on the source and target, you can disable the replication of file attributes. When attribute replication is disabled, files on the target can inherit permissions from the parent directory on the target. See the TGDisableAttributeReplication server setting in the *Reference Guide*.

Anti-virus protection

- **Carbonite Availability queue**—Exclude the Carbonite Availability queue directory on the source and target from any real-time scanning or scheduled system scans. If a queue file is deleted by a process other than Carbonite Availability, unexpected results may occur, including an auto-disconnect due to the loss of queued data. The files in the source queue directory have already been scanned (cleaned, deleted, or quarantined) in their original storage location. The files in the target queue have already been scanned (cleaned, deleted, or quarantined) on the source.
- **Target data**—Exclude the copy of the source data stored on the target from any real-time scanning or scheduled system scans. The files have already been scanned (cleaned, deleted, or quarantined) on the source. If the replicated data on the target must be scanned for viruses, configure the virus protection software on both the source and target to delete or quarantine infected files to a different directory that is not being protected. If the virus software denies access to the file because it is infected, Carbonite Availability will continually attempt to commit operations to that file until it is successful, and will not commit any other data until it can write to that file. Additionally, if the virus protection software cleans the file, an operation to clean the file will likely also be replicated from the source, which may result in file corruption.

Hardware configurations

- **NIC teaming**—If you are using NIC teaming, set it up for fault tolerance, not load balancing.
- **Device drivers**—Keep your hardware device drivers, especially NIC drivers, up-to-date.
- **Port speed and duplex**—Set static values for port speed and duplex on NICs and switches, if possible.

Full server optimizations

Review the following optimizations when you are using a full server protection job.

- **Third machine to run Carbonite Replication Console**—Ideally, you should use a third machine to run the Carbonite Replication Console and set up protection and to perform failover and reverse. If you do not use a third machine, you may need to remove and reinsert your servers (using reserved IP addresses) into the console. If you use a third machine, it must be able to communicate with the reserved IP addresses.
- **NIC configuration**—If you are planning to failover the IP address of the source, use a separate NIC and separate network for a Carbonite Availability reserved IP address that will not be failed over. If you are unable to do that and just one NIC is used for both production and reserved IP addresses, disable DNS registration on the NIC. If you are not going to failover the IP address of the source, an additional NIC and address is not necessary. In this case, Carbonite Availability will block the DNS record for that address while it is failed over.
- **Single NIC**—If you have to use only one NIC, disable DNS registration and ensure the reserved IP address is first in the list of IP addresses.
- **Disabling DNS registration**—Disabling DNS registration for the reserved IP address ensures that an end-user is not communicating to the original source when it is failed over because two different DNS records will point to two different servers.
- **Unnecessary target components**—Do not install applications, features or language packs on the target that are not essential to that machine. These can slow the failover process and/or cause mirroring issues.
- **Vendor applications**—Disable or remove any vendor applications that you are not using. These applications may cause issues after failover, like application crashes or even server crashes. These applications may not be in standard application directories, like C:\Dell or C:\cqpsystem, and those directories should be added to the list of staged folders. Also, they may use resources (like C:\INetPub), which may be needed after failover. Those directories would also have to be staged.
- **Hardware maintenance**—If you replace a NIC after you have established full server protection, you should delete and re-create your job.

Application optimizations

Review the following optimizations when you are using a SQL Server job or if you are protecting any other application on your source.

- *General applications* on page 395
- *SQL* on page 395

General applications

- **Application services on the target**—Ensure that all application services on the target are stopped and set to manual.
- **Connection mappings**—When protecting an application server, select the **One To One** mapping when creating a files and folders job.
- **Database backups**—Pause the target while you perform a backup of database files stored on the target because the database and log files must be backed up when they are at the exact same point in time. For example, say the back up of the file mydatabase.mdf begins on the target. While the backup program has access to the file, Carbonite Availability cannot write to the file. When the backup completes, Carbonite Availability writes to the file. Carbonite Availability also writes to the corresponding mydatabase.ldf file. When the backup gets to the mydatabase.ldf file, it no longer matches the .mdf file. The database may require special startup procedures, which may result in incomplete transactions being left in the database or data loss. To workaround this scenario, pause the target before starting the backup and then resume the target when the backup is complete.

SQL

- **Memory**—Typically, SQL uses all available memory. You may want to consider limiting SQL memory usage to allow the Double-Take service to function without running out of memory.
- **Temp database**—Check with your application vendor to determine if the temp database is used and needed. If it is not needed, you can exclude it from replication. For example, SQL Server re-creates the tempdb database file each time it starts, so any tempdb data that gets replicated to the target will never get used. Writes to the tempdb database may account for a significant percentage of writes to all SQL Server files, so excluding the tempdb files may result in much less replication traffic. If the database application you are using uses the temp database file (for example in Prophecy, PeopleSoft, and BizTalk) or if you are uncertain, do not exclude it from replication.
- **SQL service account**—Configure the source and target to use the same domain account to start the SQL services, if possible. This eliminates the need to move SQL Service Principal Names (SPNs) during failover and failback. If you have to use different accounts, Kerberos authentication will require the Service Principal Names to be failed over.

Chapter 13 Security

To ensure protection of your data, Carbonite Availability offers multi-level security using native operating system security features. Privileges are granted through membership in user groups defined on each machine. To gain access to a source or target, the user must provide a valid operating system user name and password and the specified user name must be a member of one of the Carbonite Availability security groups. Once a valid user name and password have been provided and the source or target has verified membership in one of the security groups, the user is granted appropriate access to the source or target and the corresponding features are enabled in the client. Access is granted on one of the following three levels.

- **Administrator Access**—All features are available for that machine.
- **Monitor Access**—Servers and statistics can be viewed, but functionality is not available.
- **No Access**—Servers appear in the clients, but no access to view the server details is available.

Although passwords are encrypted when they are stored, Carbonite security design does assume that any machine running the client application is protected from unauthorized access. If you are running the client and step away from your machine, you must protect your machine from unauthorized access.

- *Adding users to the security groups* on page 397
- *Changing the account used to run the Double-Take service on Windows servers* on page 398

Adding users to the security groups

The security groups are automatically created during the installation process. The Double-Take Admin and Double-Take Monitors groups are automatically created and the local administrator and domain administrator are automatically added to the Double-Take Admin group during installation.



If Carbonite Availability is installed on a Windows member server, it will use the local groups. If an Active Directory user is granted access to the Active Directory Double-Take Admin or Double-Take Monitors groups, the user or domain group must also be granted access to the local Carbonite Availability groups. If Carbonite Availability is installed on a Windows domain controller, the Active Directory group will provide sufficient access. The groups are created in the Users Container and need to stay here. If the groups are not there, users will be unable to log into Carbonite Availability on the domain controller.

Users that need administrator access to Carbonite Availability must be added to the Double-Take Admin group. Users that need monitor only access must be added to the Double-Take Monitors. In both cases, local users, domain users, or global groups may be added to the local groups.

See your Windows documentation for instructions on adding, deleting, or modifying users in a security group.

Changing the account used to run the Double-Take service on Windows servers

By default, the Double-Take service on Windows servers is configured to log on as the system account. If you want to select a specific account to run the service, use these instructions.



If you are protecting an entire server, you cannot modify the account used to run the Double-Take service. Otherwise, the full server protection will not function correctly.

1. Modify the user account that the Double-Take service is using.
 - a. Open the Double-Take service properties and select the **Log On** tab, select **This Account**, and enter a valid domain account.
 - b. Enter the password for this account.
 - c. Click **OK** to save these settings.
 2. Grant an additional user right to the account you are using to run the Double-Take service.
-



If domain-level policy settings are defined (through **Domain Security Policy**, **Security Settings**, **Local Policies**, **User Rights Assignment**), they will override local policy settings.

- a. Select **Local Security Policy** from Administrative Tools.
 - b. Expand the **Local Policies** folder and highlight the **User Rights Assignment** folder.
 - c. Double-click the option **Act as part of operating system** on the right pane of the screen.
 - d. Add the user that you selected to run the Double-Take service and click **OK**.
 - e. Exit the Local Security Settings dialog box. This user is now configured to run the Double-Take service.
3. Add the domain account to the local administrator group.
 - a. Select **Computer Management** from Administrative Tools.
 - b. Expand the **Local Users and Groups** folder and highlight the **Groups** folder.
 - c. Right-click on the **Administrators** group on the right pane of the screen and select **Add to Group**.
 - d. Click **Add**.
 - e. Locate the domain account that you are using for the Double-Take service. Select that account and click **OK**.
 - f. Click **OK** to close the Administrators Properties dialog box.
 - g. The domain account is now added to the local administrator group. Close the Computer Management window.