

A large, light-orange circular graphic with a stylized 'S' icon (two parallel diagonal lines) in the center, serving as a background for the main text.

SCALE COMPUTING MOVE
POWERED BY
CARBONITE MIGRATE
GETTING STARTED

Table of Contents

INTRODUCTION	2
TERMINOLOGY	2
REQUIREMENTS	3
CARBONITE MIGRATE	3
SCALE COMPUTING HYPERCORE	4
YOUR ENVIRONMENT	4
SIZING A HYPERCORE VM TARGET SERVER	5
ADDITIONAL SPACE RECOMMENDATIONS	6
INSTALL CARBONITE MIGRATE CONSOLE	6
INSTALL CARBONITE MIGRATE ON SOURCE AND TARGET WINDOWS SERVERS	7
ADD SOURCE AND TARGET SERVERS TO THE CARBONITE CONSOLE	10
LICENSING SOURCE SERVERS FOR CARBONITE JOBS	11
(PRIOR TO CARBONITE MIGRATE 8.2 ONLY) LICENSING TARGET SERVERS FOR CARBONITE JOBS	13
SETUP AND CONFIGURE A JOB	16
MONITOR JOB MIRRORING	20
JOB FAILOVER	22
RUN THE FAILOVER	23
COMMON TROUBLESHOOTING STEPS FOR CARBONITE MIGRATE	24
LICENSING	24
COLLECTING TROUBLESHOOTING LOGS FOR CARBONITE SUPPORT	25
COMMON CONSOLE ISSUES WHEN ADDING SERVERS OR CONFIGURING JOBS	25
TIPS AND TRICKS	26
RESOURCES	26
CARBONITE	26
FEEDBACK & SUPPORT	26

INTRODUCTION

Scale Computing Move Powered by Carbonite Migrate allows you to migrate a Windows or Linux physical, virtual, or cloud server workload to the Scale Computing Platform (SC//Platform) solution with real-time replication and near-zero downtime failover—all managed from a single, intuitive user console. Carbonite software is hardware agnostic and operates from within all supported operating systems on physical or virtual hosts. All steps and screenshots used in this document will utilize the Carbonite Migrate version 8.2 interface.

TERMINOLOGY

Due to the partnership inherent in Scale Computing Move Powered by Carbonite Migrate, this guide provides the intended definition for common terminology in use between Scale Computing and Carbonite.

- **Move** — Carbonite software product providing a near-zero downtime migration to the SC//Platform solution from a physical or virtual host. Scale Computing has partnered with Carbonite to offer Move as a migration solution for new Scale Computing customers.
- **Console Server** — The single management interface for all Carbonite Migrate tasks. No licensing is required to install the Carbonite Migrate Console service.
- **Source Server** — The existing server in the local environment, physical or virtual, intended to be migrated to the SC//HyperCore system. A Source server will require a single Source license. A Source license is active for 30 days once applied to a server before it expires.
- **Target** — The HyperCore VM used as a temporary destination for the existing Source environmental workload.
- **Job** — A broad term that entails the connection between a specific Source and Target server for the purpose of workload migration through the Carbonite Migrate software. This includes the data being mirrored between the Source and Target as well as the final failover from Source to Target server.
- **Queue** — A temporary storage location and buffer for replication data that cannot be immediately transmitted to the Target for any reason, such as high transmission volume, or a locked file causing processing to back up.
- **Mirroring** — The process of creating a baseline copy of all server attributes and data from the Source server to the Target server.
- **Replication** — The continuous monitoring and updating of real-time byte level changes between the Source and Target servers, keeping the attributes and data synchronized until the final failover is completed.
- **Failover** — The process of cutting over from running the server workload on the Source server to running the workload on the Target server. This includes any attributes that may have been locked by the Source server operating system until this point, shutting down the Source server, and bringing all workload access up on the Target server. The intention is for the environment to see nothing more than a “reboot,” although all workloads, connections, and machine identity components have moved from the Source server to the Target HyperCore VM.

REQUIREMENTS

Prior to using Scale Computing Move Powered by Carbonite Migrate, ensure that the following requirements are met.

CARBONITE MIGRATE

- Ensure you meet all of Carbonite Migrate requirements for use with your chosen Console, Source, and Target servers. [Refer to the Carbonite Migrate Installation & Upgrade notes](#) for more detailed requirements applicable to your Full Server migration.
- Have at least one Source server license for Scale Computing Move. Scale Computing Move requires 1 Source server license for each physical or virtual machine you wish to migrate to the SC//Platform solution. **Licenses are good for 1 year from their purchase date, whether they have been activated or not.**

NOTE

As of Carbonite Migrate Version 8.2, Target licenses are no longer necessary.

- You have 30 days to complete a Scale Computing Move job once the Source license is activated on a server.
- Download the Carbonite Installation package from the email link the Scale Computing Operations team sends you after your purchase. Ensure you install the appropriate Carbonite Migrate version for your operating system.
- Configure the Carbonite Migrate Console server on a Windows system that is accessible to the Target HyperCore system and any Source servers, and ensure the Console is separate from any servers intended as a Source or Target in a Carbonite Migrate job.

SCALE COMPUTING HYPERCORE

- Have a fully configured and initialized Scale Computing HyperCore system.
- Know how to create a HyperCore VM in the HyperCore web interface. Instructions for creating a VM are provided in the [Scale Computing HyperCore User Guide](#).

NOTE

The HyperCore VM used as a Target for the Carbonite Migrate job should be configured as a duplicate of the Source server. Ensure the OS, the number of drives, and all drive lettering within the OS of the HyperCore VM Target matches the Source server exactly.

It is not required to match the drive sizes on the HyperCore VM Target, although the size cannot be less than the existing data allocation on the Source server. See [Sizing a HyperCore VM Target Server](#) for more information on Target server sizing.

YOUR ENVIRONMENT

- Make sure all of the server names you intend to use in Carbonite Migrate jobs have been given names you can use to easily reference them in the Console.
- Ensure you have a mouse with two buttons available. Right click cannot always be mimicked on a VM console.
- Ensure all of the Source, Target, and Console servers can communicate on your network.
- An Active Directory environment is not necessary but can assist in some troubleshooting steps depending on your environment.
- The Source and Target OS type must match.

SIZING A HYPERCORE VM TARGET SERVER

As part of the Carbonite Migrate process, a copy of the Source server state data will be staged on the Target server boot volume in a folder called Staging-SSM. If the Target server's boot volume does not have enough space to accommodate the Source data and staging folder, the job will become stuck in a perpetual "Retrying" state and unable to complete synchronization. Use the steps below to predict approximately how much space will be needed for the staging folder by calculating the size of the current folder and data structure on the Source server boot volume.

This is also a great opportunity for future server planning to utilize the HyperCore thin provisioning capability. By anticipating future growth on the HyperCore VM Target server and adjusting the drive size(s) accordingly, you can accommodate additional server growth without penalty.

Calculate the combined size of the following folder structures on the boot volume of the Source server.

- Documents and Settings
- Program Files
- Program Files (x86)
- Program Data
- Windows
- Users
- Any other folder manually selected for staging

ADDITIONAL SPACE RECOMMENDATIONS

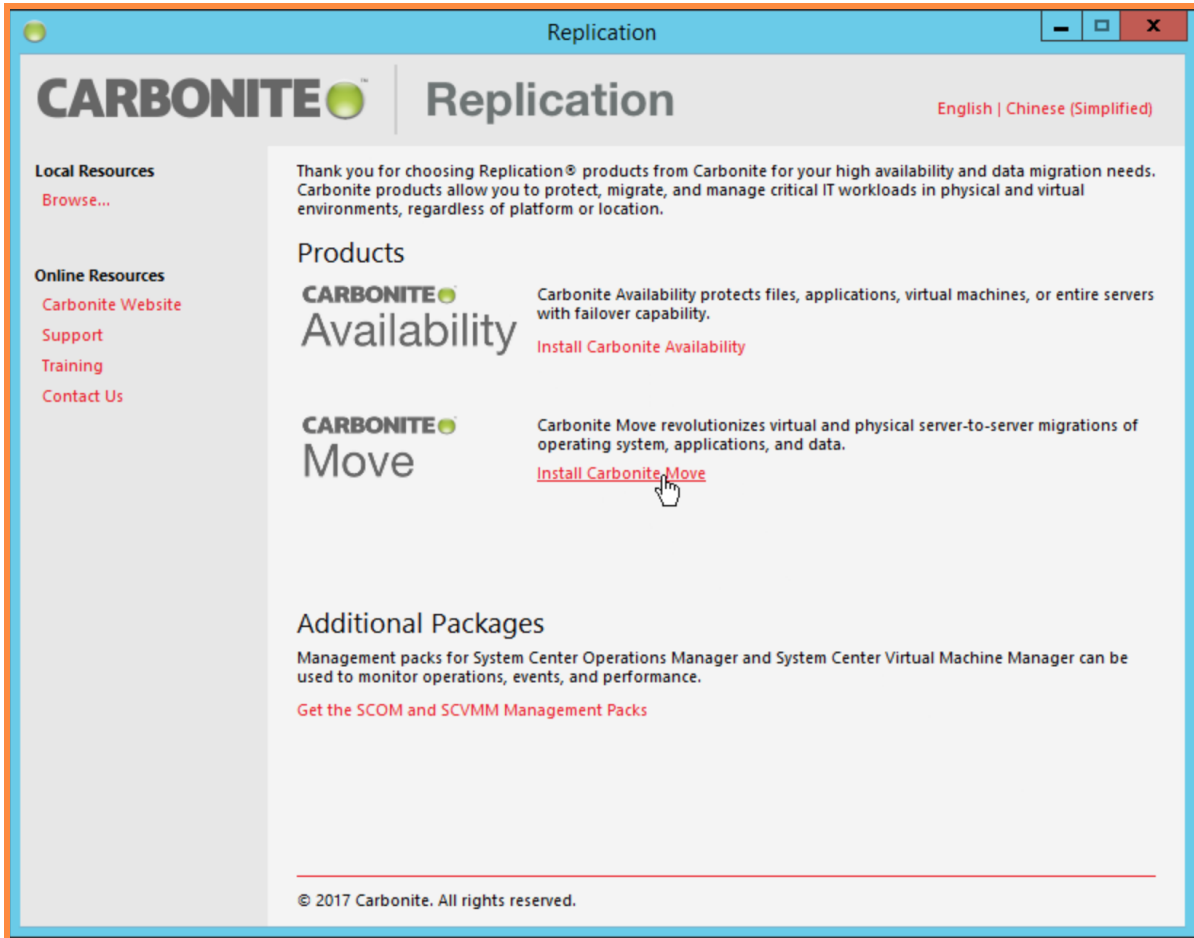
Carbonite recommends the following estimates for **additional space on the Target server boot volume beyond the calculated consumed space on the Source server boot volume** determined in Step 1. These minimums are for a clean operating system installation. Operating system customizations, installed applications, and user data will increase the disk space requirement.

- Windows 2008 — at least 7-9 GB additional
- Windows 2008 R2 — at least 10 GB additional
- Windows 2012—at least 14 GB additional
- Windows 2012 R2—at least 15 GB additional
- Windows 2016—at least 25 GB additional
- Windows 2019—at least 23 GB additional
- Windows 2022—at least 26 GB additional

INSTALL CARBONITE MIGRATE CONSOLE

Follow the steps below to install and configure the Carbonite Migrate Console server.

1. Unzip and run the Carbonite Installer on your chosen Console server. The Console should be separate from any servers planned to be used as Source or Target servers in Scale Computing Move jobs to the HyperCore system.

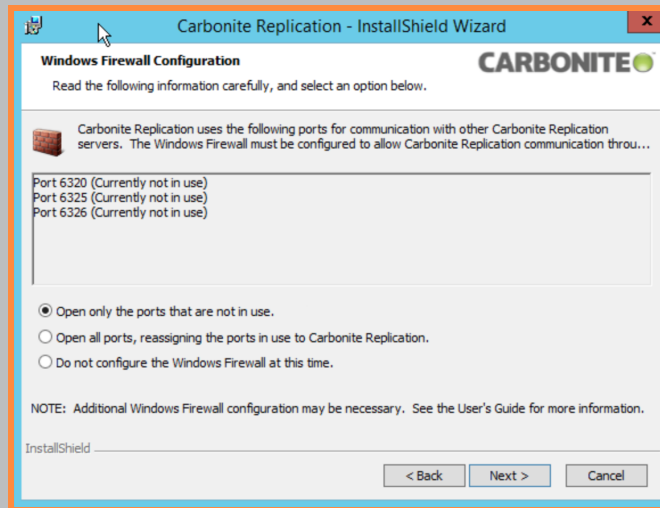


2. Choose Install Carbonite Migrate. The .NET Framework 3.5 SP 1 or later and Visual C++ 2008 SP 1 Runtime are required. The installer will ask to install or enable .NET and Visual C++ if the requirements are not met. As with typical .NET Framework and Visual C++ installations, prior versions are not overwritten; new installations are simply added to the system.
3. The installer will ask you to check for updates. If you are unsure which license(s) you have, select No, skip this step and continue the current installation. Click Next when you have made the selection.
4. Accept the License Agreement. Click Next.

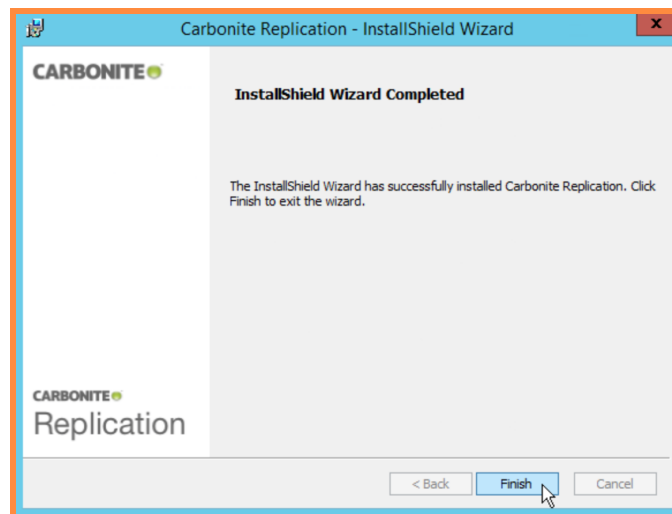
- Choose Client Components Only when asked which components you would like to install for the Console server. Click Next.

NOTE

If ports 6320, 6325, or 6326 are being blocked by a firewall, the install will generate a dialog box, shown in the image below. These ports must be unblocked for Carbonite Migrate to work correctly. If using Windows Firewall, select the appropriate option and click Next. If you are using a third party firewall, manually unblock the ports before you click Next to proceed.



- Click Install to complete the Console installation.
- Click Finish when the InstallShield Wizard notifies you of the completed installation. The Carbonite Migrate Console components have been installed on the server



INSTALL CARBONITE MIGRATE ON SOURCE AND TARGET WINDOWS SERVERS

Follow these steps to install Carbonite Migrate on the Source and Target servers.

NOTE

Any Source Server will require Server Components (Server Components Only in the Setup Type screen during installation).

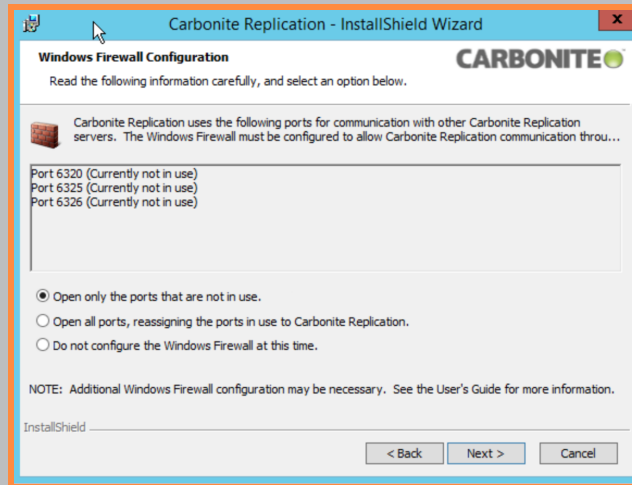
Any Target Server will only require Server Components (Server Components in the Setup Type screen during installation).

An intended Source or Target server should never double as the Carbonite Migrate Console. A separate Windows-based Console server should be configured to be accessible on the network.

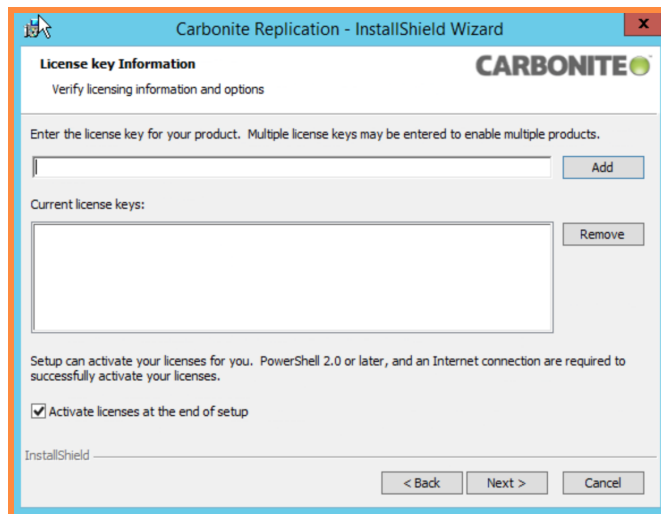
1. Unzip and run the Carbonite Installer on your chosen Source or Target server.
2. Choose Install Carbonite Migrate. The .NET Framework 3.5 SP 1 or later and Visual C++ 2008 SP 1 Runtime are required. The installer will ask to install or enable .NET and Visual C++ if the requirements are not met. As with typical .NET Framework and Visual C++ installations, prior versions are not overwritten; new installations are simply added to the system.
3. The installer will ask you to check for updates. If you are unsure which license(s) you have, select No, skip this step and continue the current installation. Click Next when you have made the selection.
4. Accept the License Agreement. Click Next.
5. Choose the appropriate Components listed in Step 1 when prompted. Click Next.

NOTE

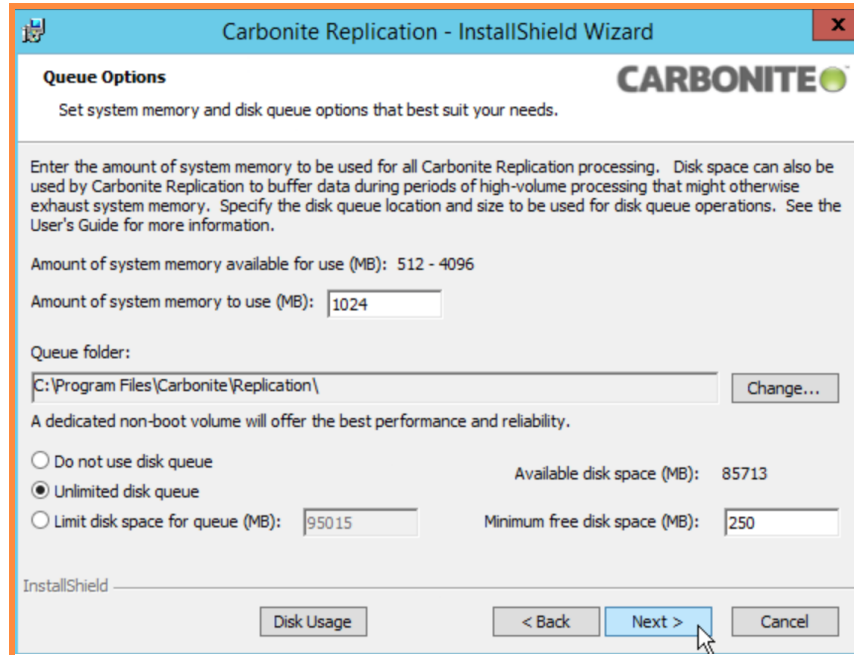
If ports 6320, 6325, or 6326 are being blocked by a firewall, the install will generate a dialog box, shown in the image below. These ports must be unblocked for Carbonite Migrate to work correctly. If using Windows Firewall, select the appropriate option and click Next. If you are using a third party firewall, manually unblock the ports before you click Next to proceed.



6. The License Key Information screen may appear and can be dismissed. All licenses for Source and Target servers will be configured later through the Console interface. Click Next.



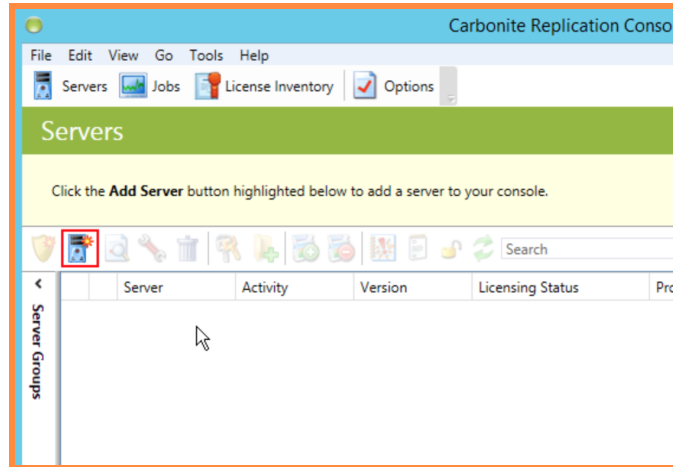
- The installer will ask how much memory to allocate to the Scale Computing Move agent, as well as the desired parameters for a disk queue to buffer data during times of high volume processing. Scale Computing recommends selecting the default values for all items. Click Next.



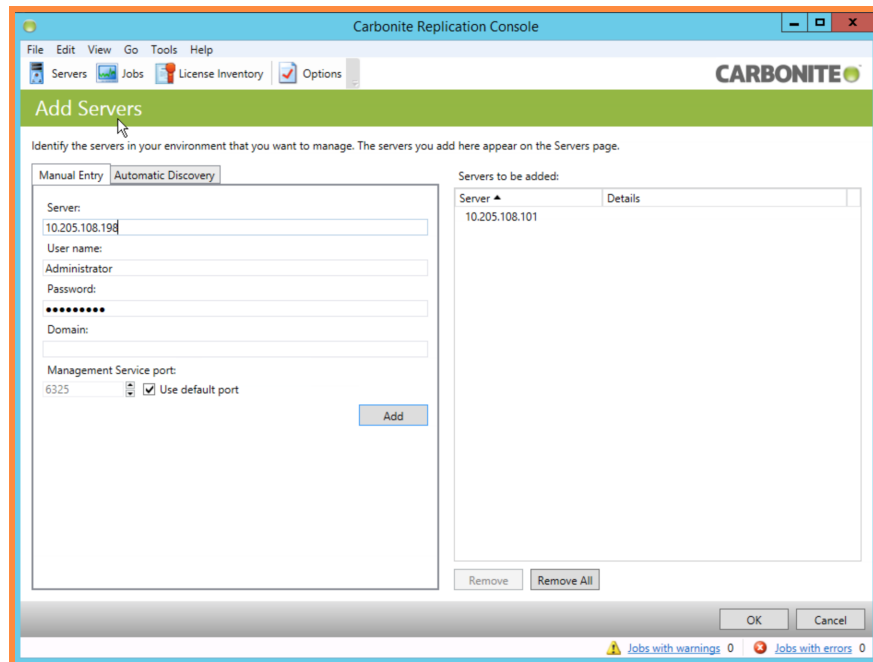
- The installer will display information on Windows Security Groups being added to the server. Click Next.
- Click Install to complete the Source or Target server installation.
- Click Finish when the InstallShield Wizard notifies you of the completed installation. The Carbonite Migrate Source and Target components have been installed on the server.
- Repeat this process for all Source and Target machines being migrated using Scale Computing Move.

ADD SOURCE AND TARGET SERVERS TO THE CARBONITE CONSOLE

1. Open the Carbonite Migrate Console on your chosen Console machine.
2. Click the Add Servers icon from the task bar in the console.



3. Select the Manual Entry tab. Although servers can be automatically added, Scale Computing recommends manually entering each to ensure the proper credentials are used for each server.



4. Click OK to add all servers listed on the right hand side.

LICENSING SOURCE SERVERS FOR CARBONITE JOBS

Source servers should be the physical or virtual machines in the existing environment intended to be migrated to the destination HyperCore VMs.

WARNING

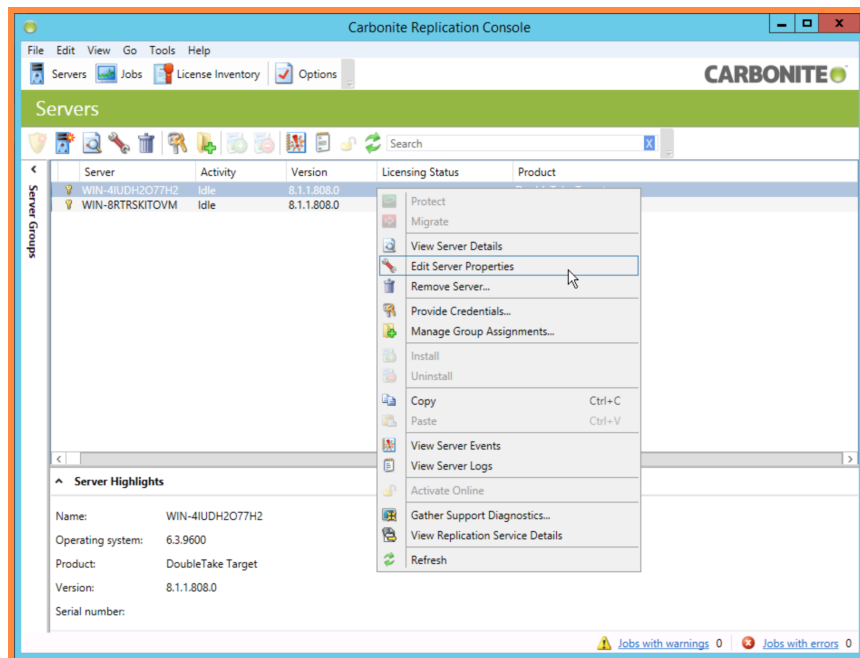
DO NOT ACTIVATE THE SOURCE LICENSE UNTIL YOU ARE READY TO COMPLETE THE MIGRATION TO SC//PLATFORM.

- Carbonite Migrate licenses are **not** active until they are applied to a Source or Target server.
 - **Target** server licenses are no longer used in Carbonite Migrate Version 8.2 and up.
- An activated Source license expires after 30 days.
- An activated Target license never expires.
 - **Target** server licenses are no longer used in Carbonite Migrate Version 8.2 and up.
- All licenses are good for 1 year from their date of purchase before they expire, activated or not.

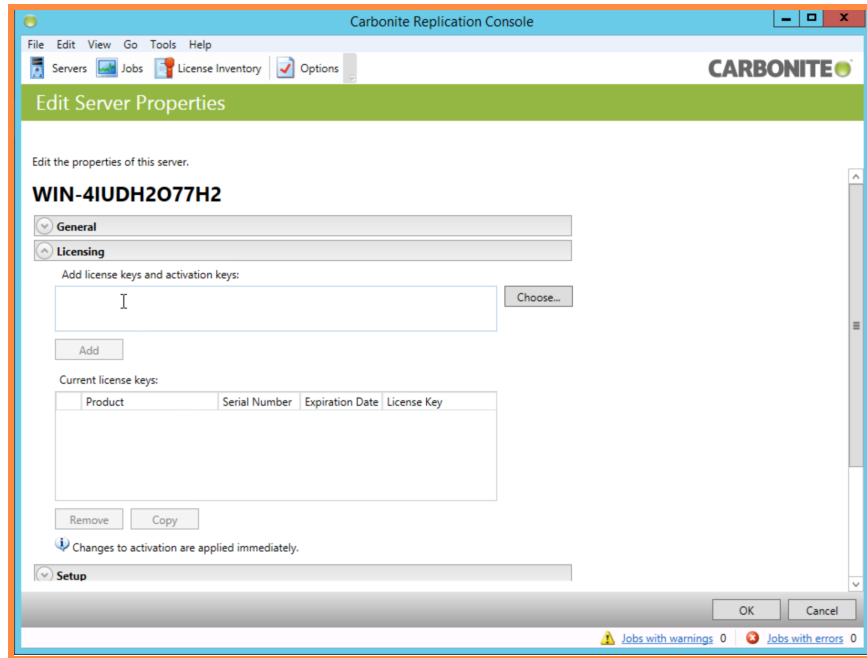
The 30 day Source license expiration period applies to individual servers. For example, if 15 keys are purchased and 1 is activated on a Source server, the single Source server license will begin the 30 day activation period on that server and the remaining 14 Source licenses will not be affected.

Failure to complete the migration (outside of technical or non-technical reasons where you have engaged with ScaleCare Support) within 30 days of activation or within 1 year of license purchase, activated or not, will require the purchase of an additional license.

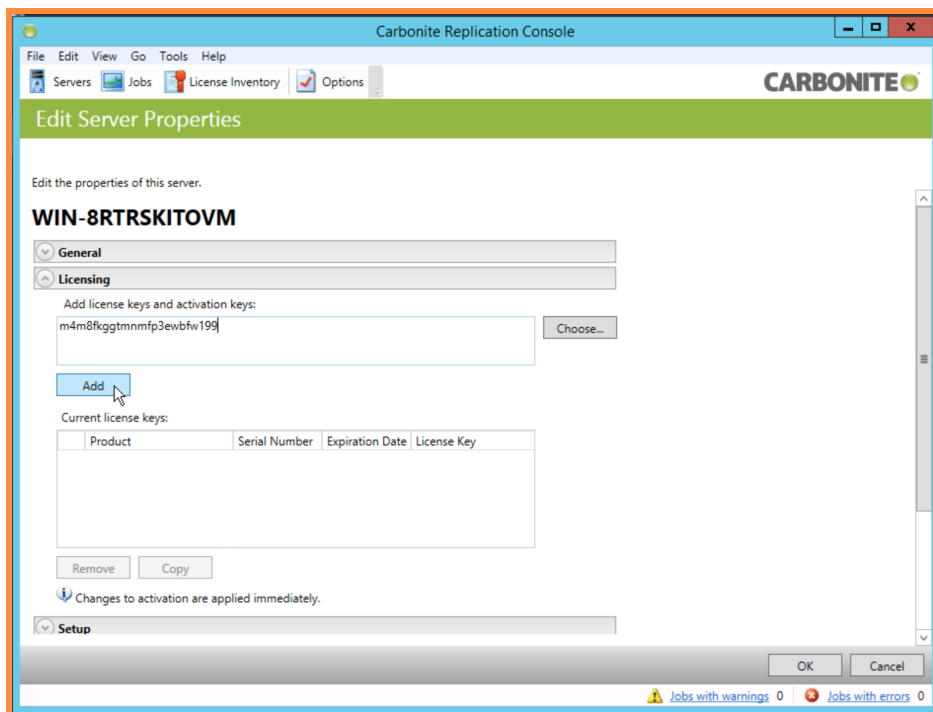
1. Open the Carbonite Migrate Console. Select the server you are going to license from the Servers view in the Console. Right-click on the selection. Select Edit Server Properties.



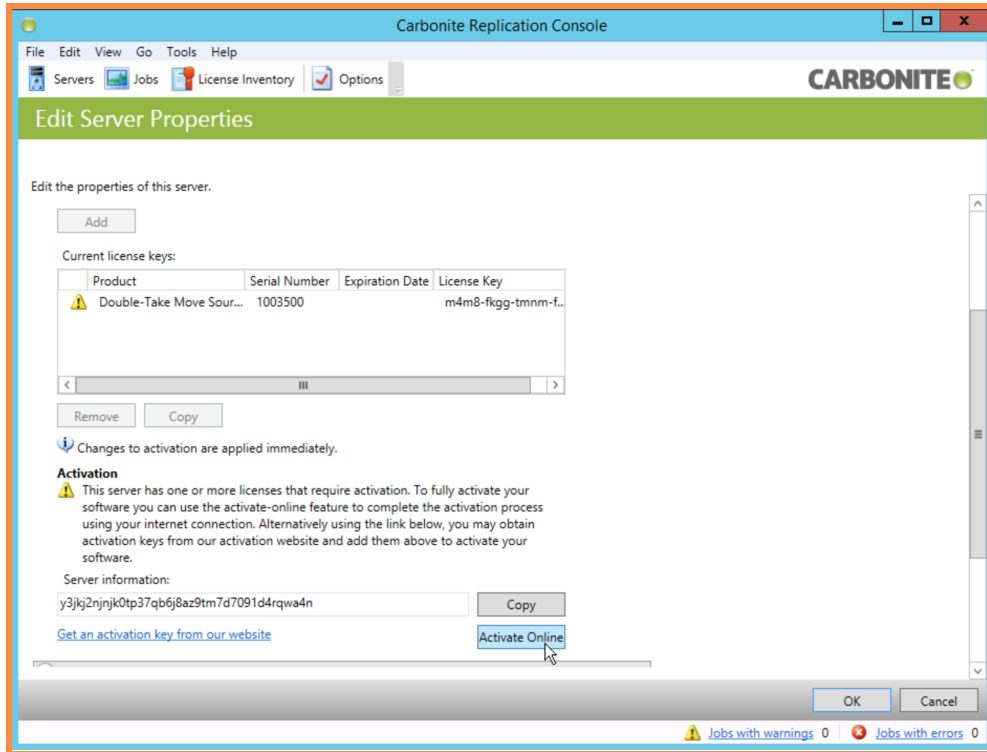
- Expand the Licensing section.



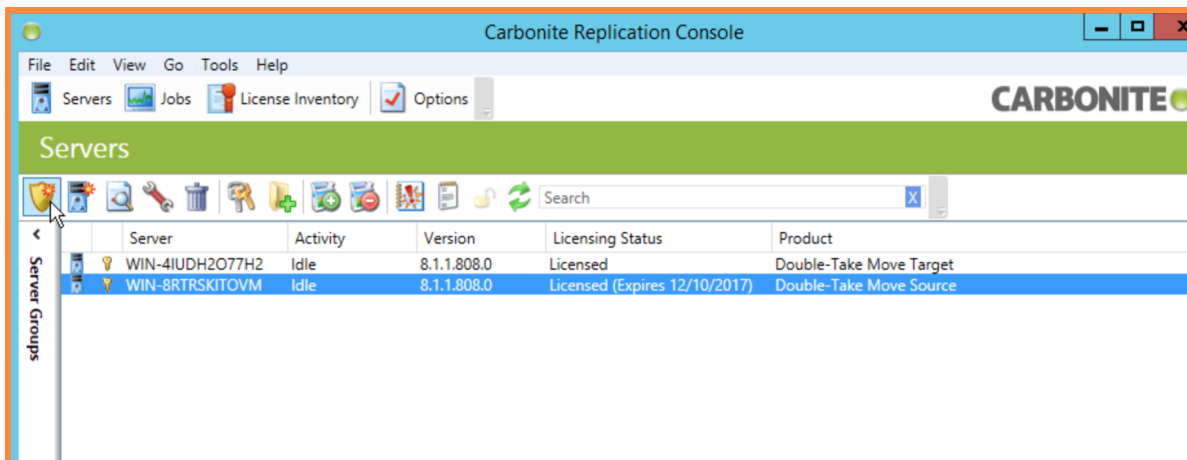
- Find the Source server license codes received in the email from Scale Computing. Enter the Source server license code into the box and click Add.



- If the license does not automatically activate, it will need to be manually activated. Scroll to the bottom of the Licensing section and click Activate Online. This will communicate with Carbonite to manually activate the license. Click OK in the bottom right corner to finish the Source server licensing process.



- In the Servers screen, you will notice under the Licensing Status that the key has been applied. Take careful note of the expiration date for the Source server license. The final failover and migration must be completed before this date.



(PRIOR TO CARBONITE MIGRATE 8.2 ONLY) LICENSING TARGET SERVERS FOR CARBONITE JOBS

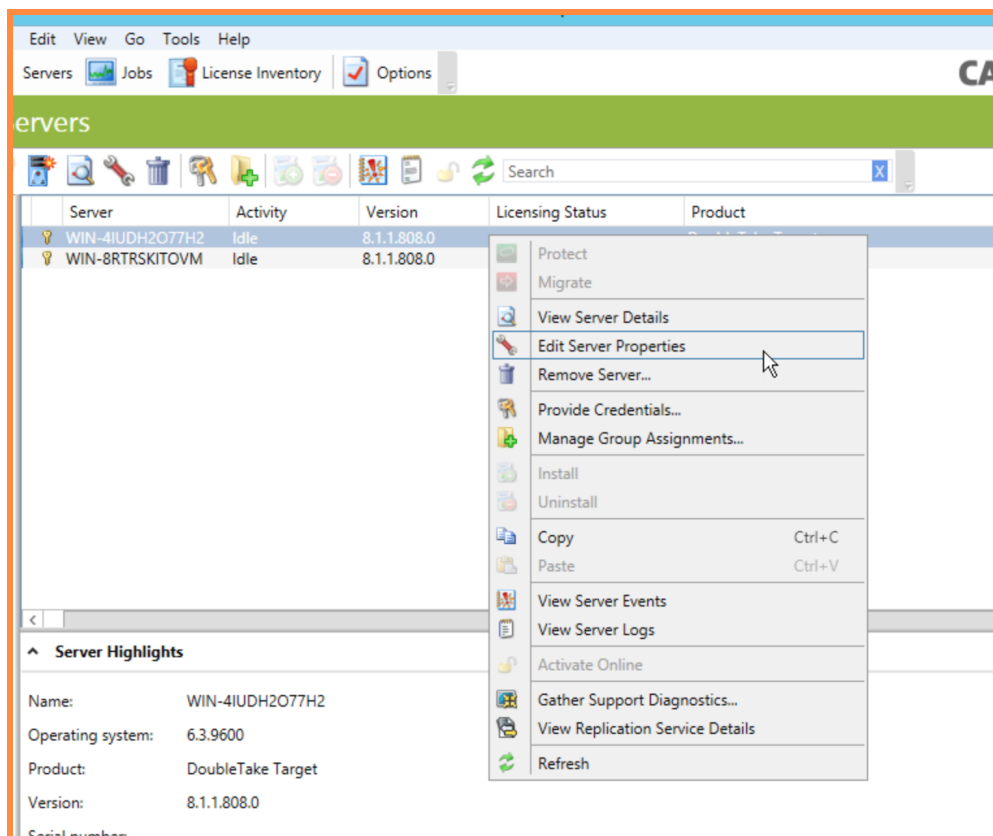
NOTE

As of Carbonite Migrate Version 8.2, Target licenses are no longer necessary. Only a Source license will be sent and will need to be activated to proceed with the Carbonite Migrate job.

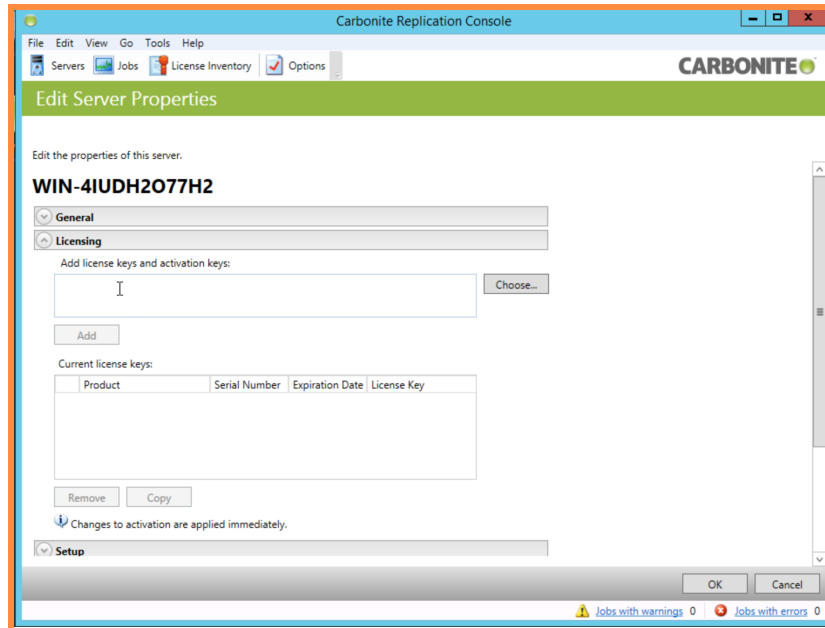
This section is only necessary if you are using a Carbonite Migrate version prior to 8.2.

Target servers should be the destination HyperCore VMs for the Source server data (physical or virtual machines in the existing environment).

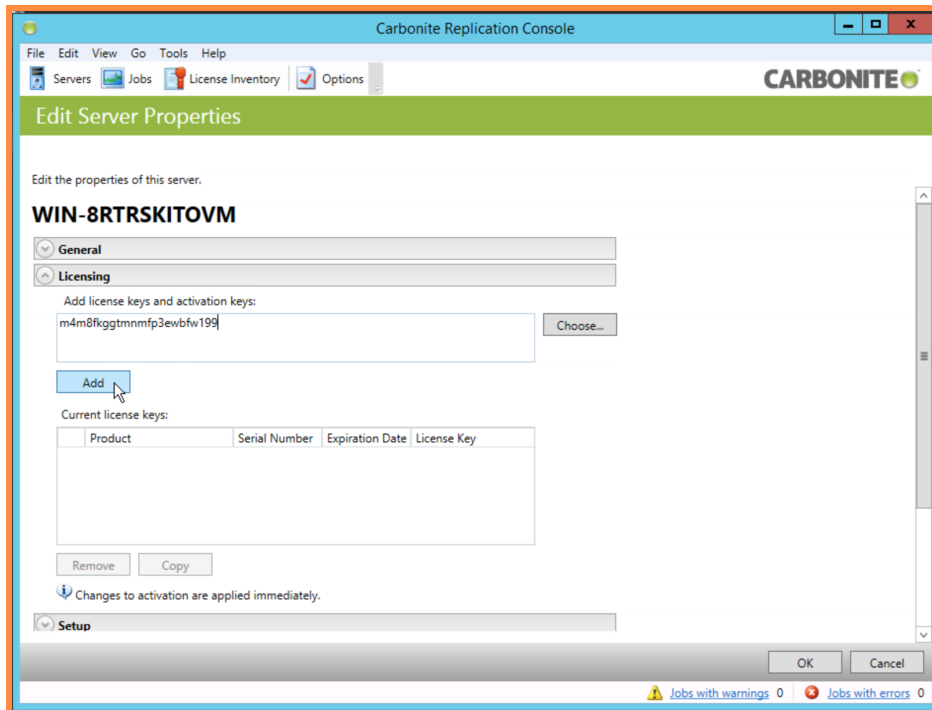
1. Open the Carbonite Migrate Console. Select the server you are going to license from the Servers view in the Console. Right-click on the selection. Select Edit Server Properties.



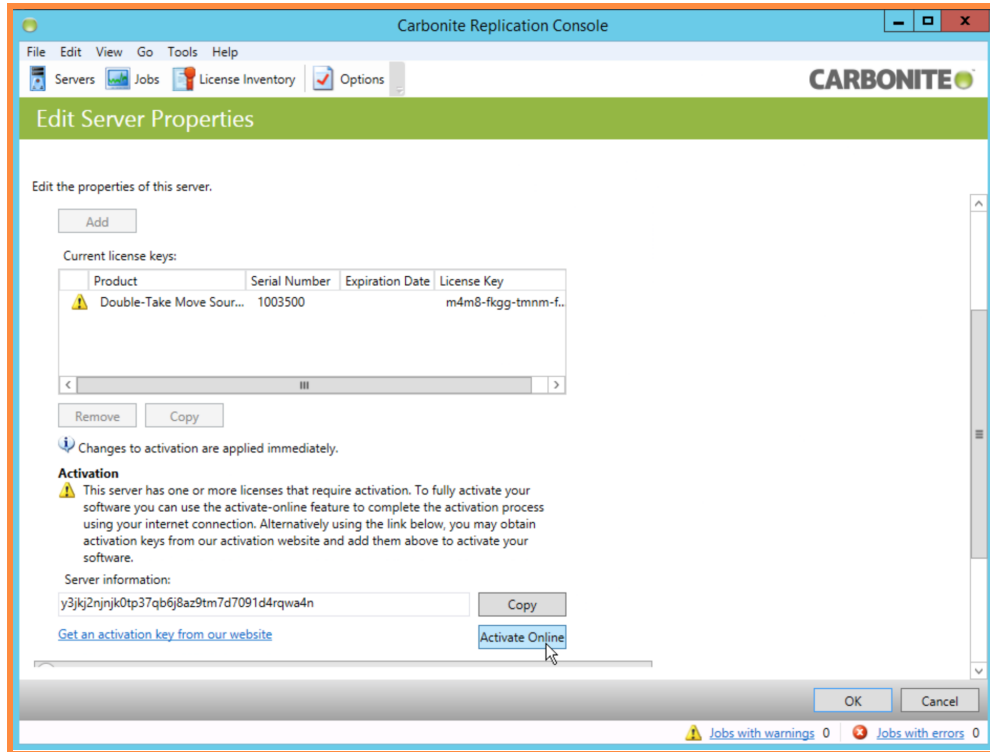
- Expand the Licensing section.



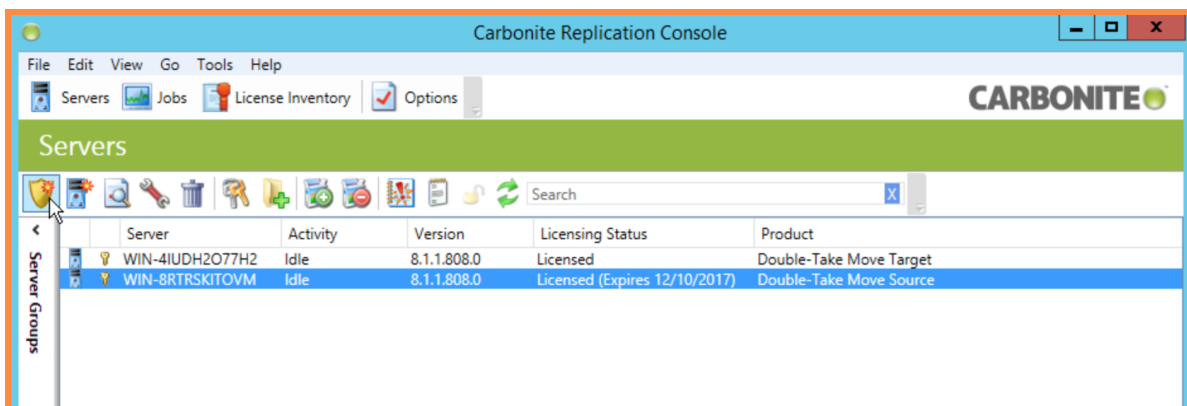
- Find the Target server license codes you received in the email from Scale Computing. Enter the Target server code into the box and click Add.



- If the license does not automatically activate it will need to be manually activated. Scroll to the bottom of the Licensing section and click Activate Online. This will communicate with Carbonite to manually activate the license. Click OK in the bottom right corner to finish the Target server licensing process.



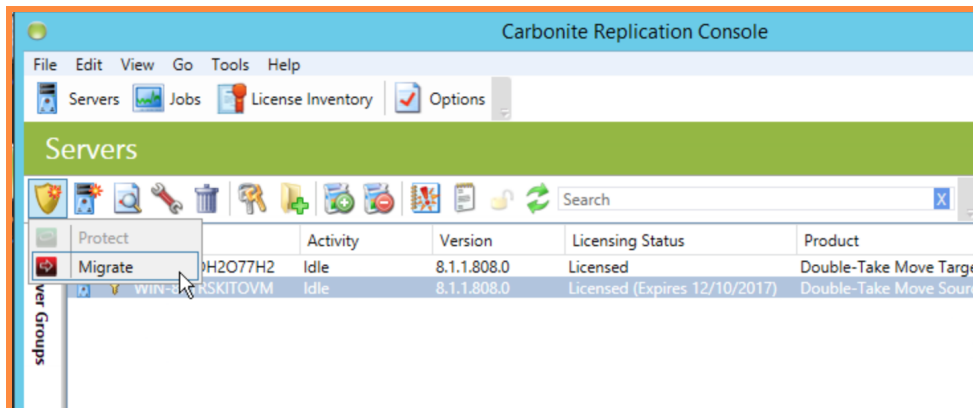
- In the Servers screen, you will notice under the Licensing Status that the key has been applied. A Target server license will not have an expiration date.



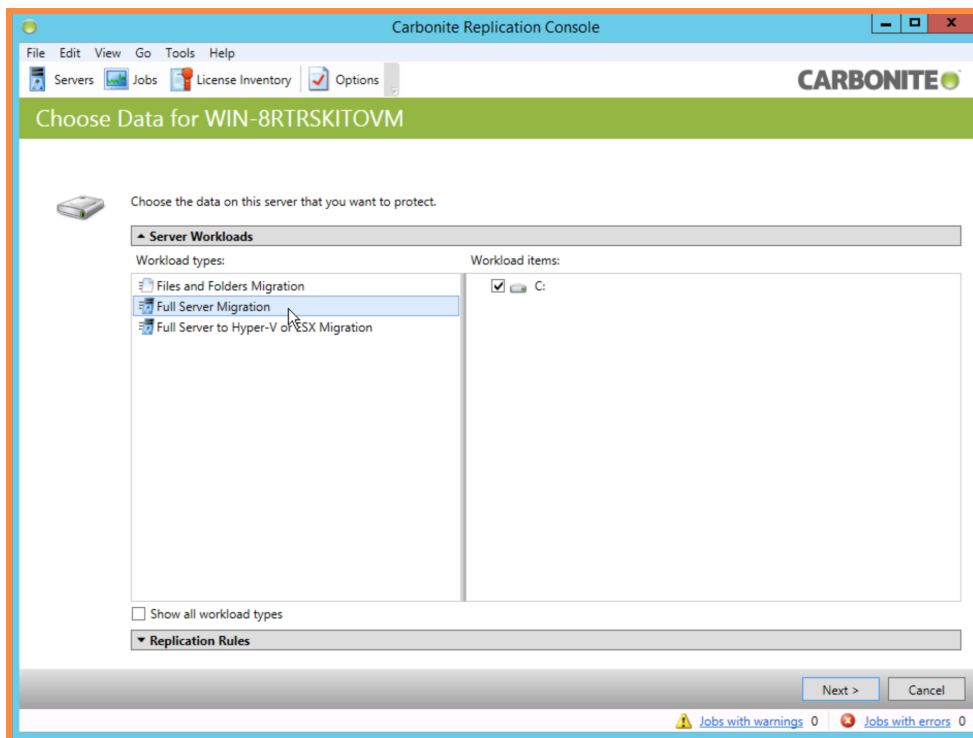
SETUP AND CONFIGURE A JOB

Once a Source and Target server have been established, a job can be set up and configured. Before configuring the job, Scale Computing highly recommends restarting the Source server to avoid issues that may arise as the server is rebooted during the job, such as interference from a pending Windows Update.

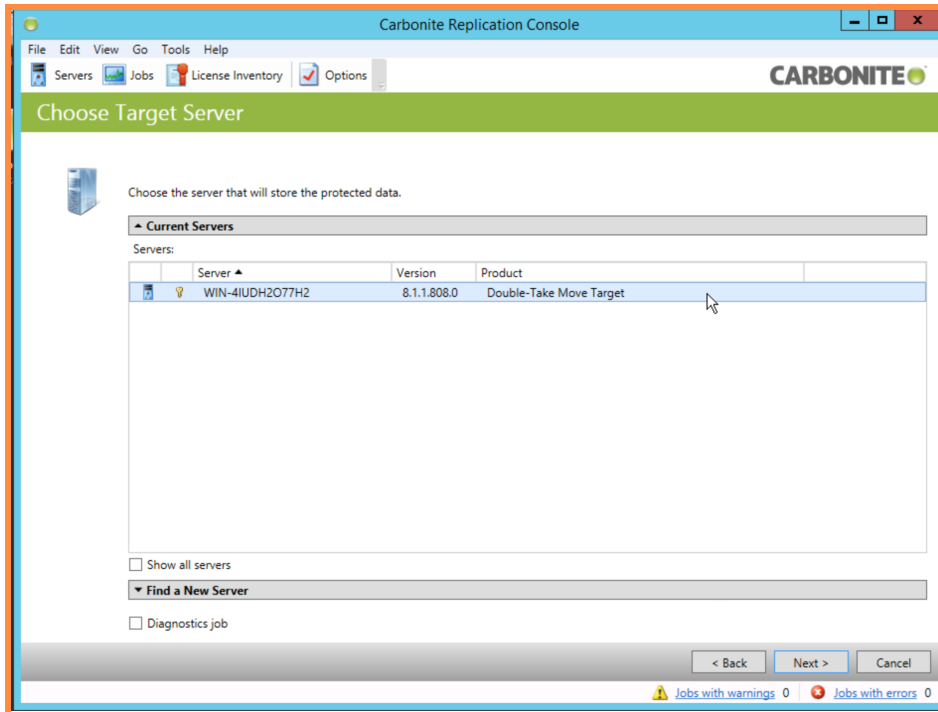
1. Click the Shield icon in the left-most slot of the task bar. It will be labeled Get Started. Click Migrate in the dropdown menu.



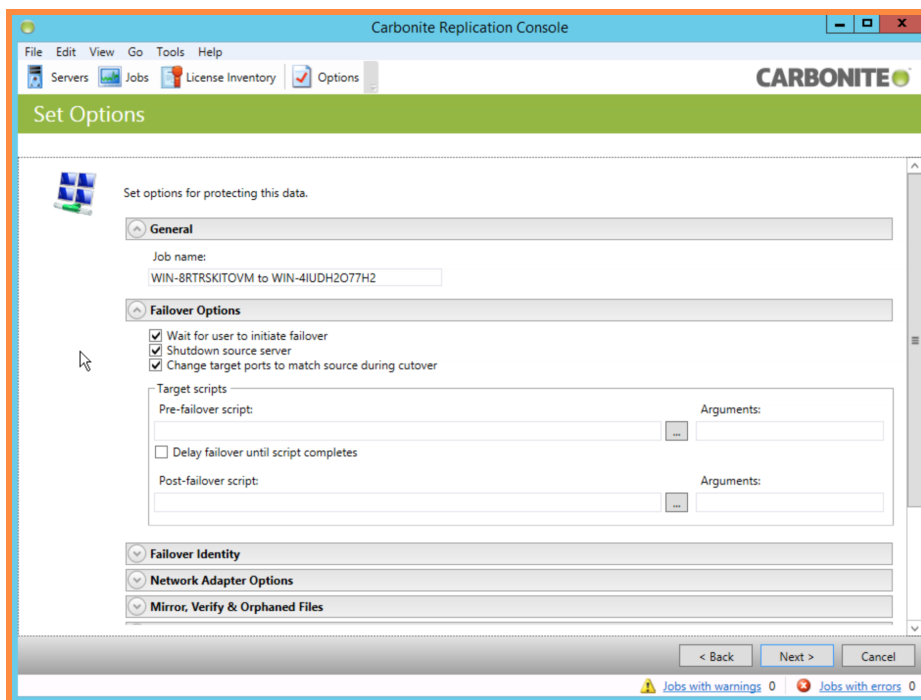
2. Select Full Server Migration. Select all of the drives on the right side you wish to migrate. Click Next.



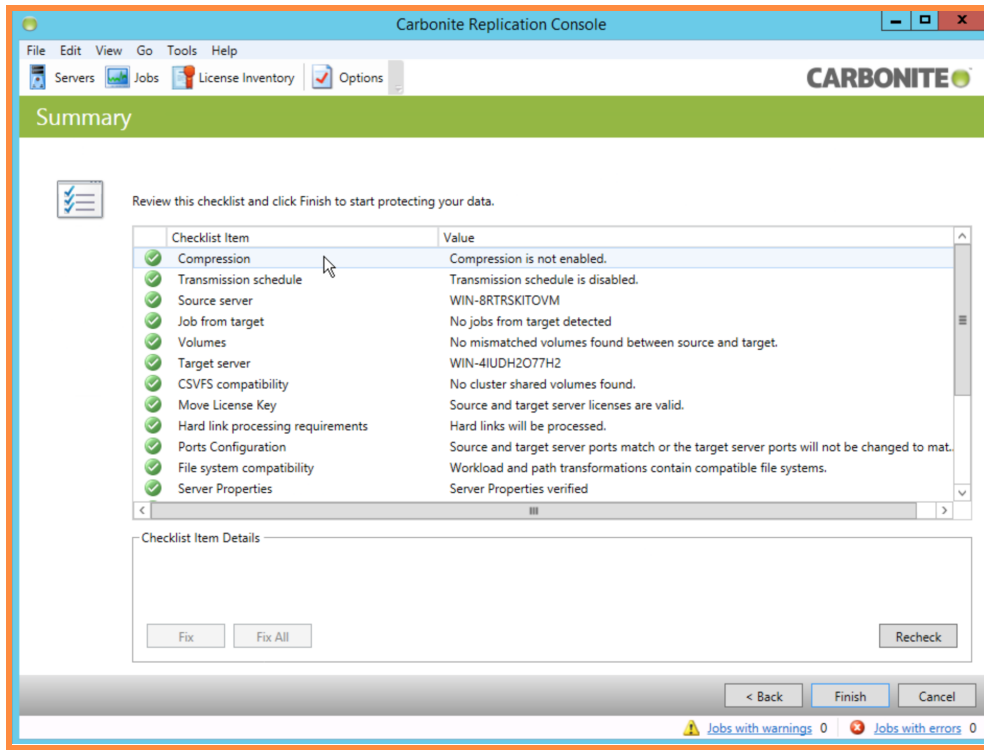
3. Select the Target Server for the job. Click Next.



4. Ensure the 3 boxes under Failover Options are selected. Scale Computing recommends using all of the default options for the Scale Computing Move Powered by Carbonite Migrate job. [Refer to the appropriate Carbonite Migrate User Guide for advanced settings.](#) Click Next.



5. Wait for the checklist to complete.
 - a. Green circles with white check marks indicate the correct configuration.
 - b. Yellow triangles with black exclamation marks are informational warnings and will not stop the job from being created, but could potentially lead to issues during failover.
 - c. Red circles with the white 'X' are critical errors that will prevent the job from being created. Any critical errors must be resolved before proceeding.



6. Click Finish to save the job settings and initiate the Carbonite Migrate job.

NOTE

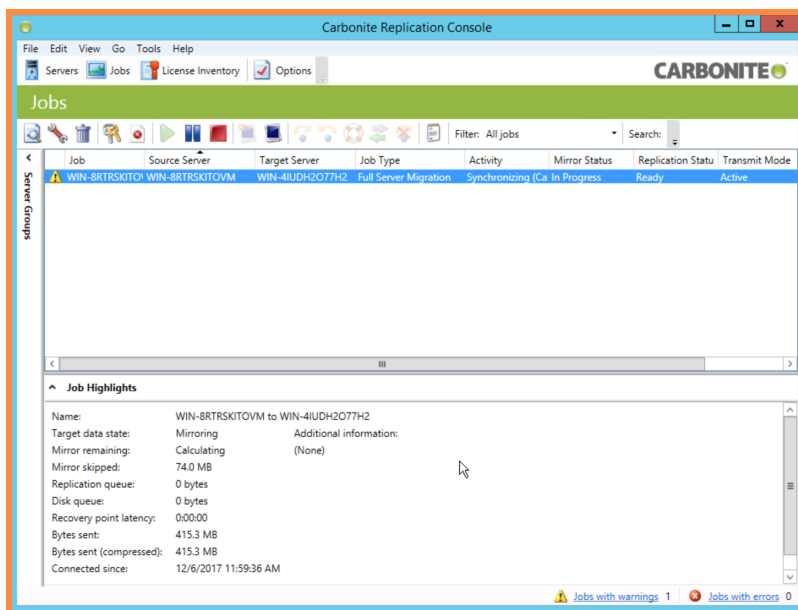
Creating the job does not initiate the failover from Source to Target. Creating the job will only begin the data mirroring between the Source and Target servers.

MONITOR JOB MIRRORING

After a job has been successfully created you will be returned to the Console Jobs screen. The new job will be shown in this view. All newly created jobs should be shown in the mirroring process with the Activity set to Synchronizing (Calculating) or Synchronizing, after calculation has completed.

The Job Highlights box below a selected job displays the detailed status of the mirroring progress. While a job is still mirroring (syncing the baseline data from Source to Target), the Target data state in the Job Highlights box will say Mirroring. Once a job has completed Mirroring to the Target server while the Source server is still running (operating system files and heavy application files may remain locked until failover), the job Activity in the Jobs screen will change to “Protecting” and the Target data state in the Job Highlights box will indicate “OK”.

Carbonite Migrate provides some additional information around mirroring progress in the Job Highlights box for each job. This includes items such as “Mirror remaining” (total number of mirror bytes that are remaining to be sent from the source to the target), “Mirror skipped” (total number of bytes that have been skipped when performing a difference mirror), and additional details on queue depths and the potential latency between recovery points.



JOB FAILOVER

The mirroring for a job must be complete before the failover can be initiated. The job Activity will show Protecting when the mirroring is complete.

The job failover is the final step in migrating a Source server workload to the HyperCore system VM Target. A Carbonite Migrate failover typically takes 15 to 20 minutes from the point the failover is initiated to the time the HyperCore VM is the primary server.

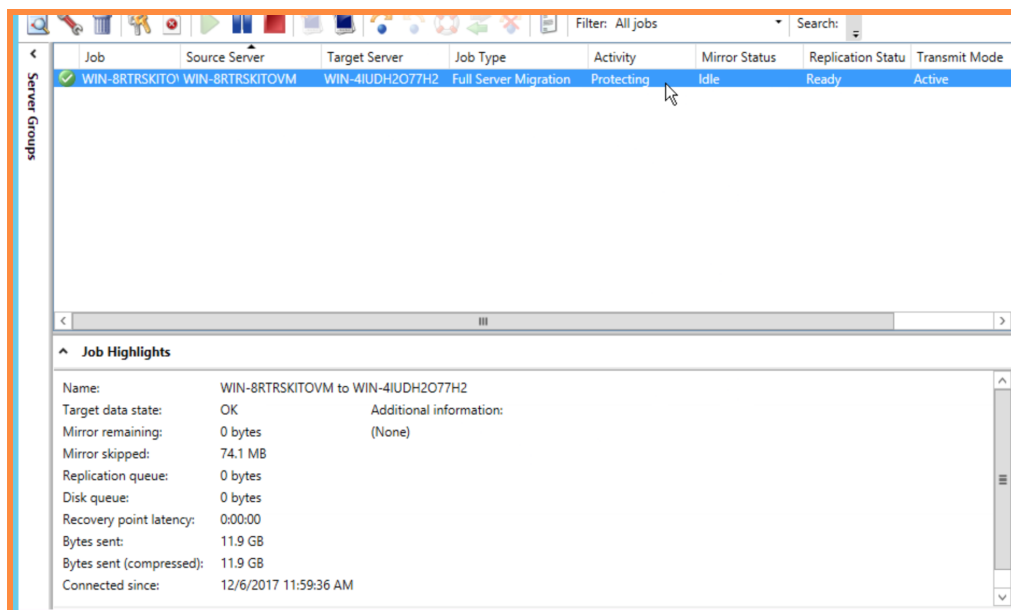
Ensure you have adequately prepared for a failover before beginning the process. Stop all active applications, and validate that the replication queue is empty (no pending transmissions). As with any other server change, a maintenance window should be established for the Carbonite Migrate failover. The server data will be unavailable for the duration of the failover.

When a failover is initiated, the source may be automatically shut down if it is still running, depending on your job configuration. The Target will stand in for the Source by rebooting and applying the Source identity, including its system state, on the target. After the reboot, the target becomes the source, and the target no longer exists.

At this point the workload will be running on the HyperCore VM Target as an exact copy of the original Source server.

RUN THE FAILOVER

1. Select the job.
2. Click Failover in the tools bar to start the Scale Computing Move Powered by Carbonite Migrate failover to the HyperCore VM Target.



3. It is possible to monitor the failover progress for the majority of the process through the Console. Source server status will be reported completely through shutdown.

NOTE

At this point the Source machine is no longer accessible. Trying to access the Source server while the HyperCore VM Target is running will cause duplicate information on the network and in Active Directory. Do not power on the Source server again without taking the necessary precautions.

4. Target server status will be accurate through the initial reboot. Due to the Carbonite service disruption for communication to the Console during the reboot, the Carbonite Console will no longer accurately report the Target status.
5. Scale Computing recommends accessing the HyperCore VM Target through the HyperCore web interface VM console to accurately monitor the remaining progress of the failover.
6. Once the failover is complete and migration is considered successful, the Carbonite Migrate license for the Source machine is consumed and **all Carbonite software and services must be uninstalled from the HyperCore VM Target.**

COMMON TROUBLESHOOTING STEPS FOR CARBONITE MIGRATE

LICENSING

Carbonite offers licenses based on the number of Source servers you wish to use. This means each Source server license code is uniquely generated with a set limit of available uses for each customer. Each individual license code instance expires 30 days from the day it is activated on the Source server. For Carbonite Migrate versions prior to 8.2, the Target license code is based on the version of Carbonite Migrate in use and does not expire.

Once a Source Carbonite Migrate license (and/or Target Carbonite Migrate license in versions prior to 8.2) has been activated on a server, it is possible to remove and reclaim the license. **Do not delete the license from the Console or delete the server itself in an attempt to reclaim a license;** you will invalidate the license and be required to purchase a new one.

Instead, from the same view where the license was activated in the Console, you will need to deactivate the license directly. If you need assistance with license deactivation and reclamation, contact ScaleCare Support.

When a Source and Target server are configured in a Carbonite Migrate job in the Console they become linked by a unique value in the Carbonite Migrate software. Any changes to the server attributes after the job is running (renaming the server, cloning or deleting the HyperCore VM Target) could result in orphaning this unique value and require ScaleCare and Carbonite Support assistance to resolve the issue.

COLLECTING TROUBLESHOOTING LOGS FOR CARBONITE SUPPORT

If you are having any issues with your Carbonite Migrate software it is possible to collect information logs to help troubleshoot your issue. The log collection program will need to be manually run on both the Source and Target server.

1. Browse to your Carbonite installation directory. Depending on what version of Move you are using the file path and nomenclature may vary. This includes Vision, Double-Take, or Carbonite for naming in the path and in the tool. Some common path examples include:
 - C:\Program Files\Vision Solutions\Double-Take
 - C:\Program Files\Double Take
 - C:\Program Files\Double-Take Software\Double-Take
2. Run the program called `DTinfo.exe`.
3. The program will collect all the relevant logs and diagnostic information needed to investigate any issue(s) further with Carbonite Support. The collection process typically takes 5 to 15 minutes.
4. The program will create a folder called Support within the installation directory when it is finished.
5. There will be a zip folder inside the Support directory containing all of the relevant information for the server.
6. Clearly rename the zip folder with the appropriate Source or Target server designation.
7. Submit both the Source and Target server zip folders to ScaleCare Support. ScaleCare Support will engage Carbonite Support on your behalf.

COMMON CONSOLE ISSUES WHEN ADDING SERVERS OR CONFIGURING JOBS

The Carbonite software is highly reliant on the communication passed between the Carbonite Console, Source, and Target background services running on the operating system. Any service or network interruption can throw off the Console view with an error indicating an inability to contact or connect to the “management service” depending on the version of Carbonite in use.

For any issue related to contacting the management service, it is often as easy as ensuring the Carbonite and/or DoubleTake services are running in the background on the specific Source or Target server. If there are still issues with the Console status and it has been verified that the Carbonite services are running on the Source and Target server, try removing and re-adding the Source or Target server to the Console. This will reset the Console connection and can often correct the issue.

TIPS AND TRICKS

Here are a few tips and tricks ScaleCare Support has found useful in our experience with Scale Computing Move.

1. Ensure both the Source and Target servers are in the same Active Directory domain if Active Directory is in use in the environment (no forested domains).
2. When configuring a Source and Target server, use the local administrator credentials for access if possible.
3. Create a small, 50 GB HyperCore VM on the Scale Computing system for easy setup, configuration, and management as a Carbonite Migrate Console.
4. The Carbonite Console has a license inventory management feature if you are confident in using the Carbonite Console and would like more flexibility in designating Source and Target licenses. Review the appropriate Carbonite User Guide for your operating system and version for more details.
5. Pro tip from the **Carbonite Migrate for Windows User Guide, Chapter 4**: “You can email Carbonite Migrate event messages to specific addresses using an SMTP mail server. The subject of the e-mail will contain an optional prefix, the server name where the message was logged, the message ID, and the severity level (information, warning, or error). The text of the event message will be displayed in the body of the e-mail message.” See the complete details from the User Guide in the Resources section.

RESOURCES

CARBONITE

- [Carbonite Documentation](#)

FEEDBACK & SUPPORT

DOCUMENT FEEDBACK

Scale Computing welcomes your suggestions for improving our documentation. Please send your feedback to documentation@scalecomputing.com.

TECHNICAL SUPPORT AND RESOURCES

There are many technical support resources available for use. Access this document, and many others, at <http://www.scalecomputing.com/support/login/>.

- [Partner Portal - Partner and Distributor use only.](#)
- [User Community - Customer focused, including our online Forum.](#)

Online Support

You can submit support cases and view account information online through the Scale Computing Customer and Partner Portals at <http://www.scalecomputing.com/support/login/>. You can also Live Chat with support through www.scalecomputing.com during standard hours Monday-Friday from 8AM - 8 PM EST.

Telephone Support

Support is available for critical issues 24/7 by phone at +1 877-SCALE-59 (+1 877-722-5359) in the US and at +44 (0) 808 234 0699 in Europe. Telephone support is recommended for the fastest response on priority issues, and the only response after standard Support hours.